THE MODEL VTU02

Is a real-time, 1/3 octave audio analyzer designed to directly connect to the TRS-80 LEVEL II or DISK computer system. No modification of the computer is necessary, and no power is required from the system. Connection of the Analyzer allows extremely rapid and versatile analysis of virtually any audio signal.

VTU02
REAL-TIME ANALYZER FOR TRS-80

Additional capabilities include voice recognition and statistical processing. The presence of the Analyzer does not prevent or affect normal operation of the computer.

********** LIMITED WARRANTY ********

The real-time analyzer is warranted against defects in materials and workmanship for a period of six months from date of purchase from Eventide or an authorized dealer.

This warranty DOES NOT COVER: Shipping damage, damage caused by improper installation, damage caused by unauthorized modification, or damage caused by physical or electrical abuse. Determination of cause of damage is in the sole province of the manufacturer. The sole liability is for repair or replacement of defective parts, and neither the manufacturer nor his agents will be liable for consequential damage.

In case of difficulty, call Eventide at (212) 581-1424 for assistance, and for a Return Authorization if it is necessary to return the unit. Any returned unit must be accompanied by a complete, detailed report of the problem, and a copy of the sales receipt to establish warranty status.

EVENTIDE CLOCKWORKS INC., 265 WEST 54TH STREET NEW YORK NY 10019

Copyright Eventide Clockworks Inc. 1979

INTRODUCTION

The Eventide Real-Time Spectrum Analyzer (Model VTU02) is designed to be connected to and operated with the Tandy/Radio Shack TRS-80 computer. After it is installed in the TRS-80, the distinction between the Analyzer itself and the computer, its keyboard, and circuitry will be somewhat nebulous (and irrelevant) in this instruction manual. Although the analyzer has no function apart from its host, the TRS-80 itself may continue to be used as it always has been. The Analyzer has no effect on the TRS-80 operation unless it is specifically called upon in the software routines to be described below.

ORGANIZATION

This manual is organized in a manner which will permit the owner to get started using the unit as soon as possible. Appendix 1 contains installation instructions, which should be followed immediately upon receipt of the unit. Appendix 2 contains test and calibration routines which may be used to assure oneself that the unit is functioning properly. Assuming the unit is installed and functioning, we begin with the software routines, followed by Theory of Operation, Applications, and technical data.

Before starting, we would like to make the following suggestions:
1: FOLLOW THE INSTALLATION INSTRUCTIONS CAREFULLY! It is not impossible to damage either the analyzer or the TRS-80 by clumsing around. It isn't easy, but it is possible.
2: Read the manual carefully. It may save you much puzzling, confusion, and 'phone bills. We have done our best to give an accurate and complete presentation. If you find any errors or have any suggestions we would be pleased to hear from you.
3: BE SURE TO RETURN THE WARRANTY CARD! This is the ONLY way we can be sure that you will receive software updates and change information. It is our intention to offer special-purpose software for the Analyzer (both in BASIC and on ROM), and we are anxious for Analyzer owners to know what is available.

We are glad you have selected this product and will be happy to assist you in its proper application and with any problems you may encounter. We have a special 'phone number for our computer-related products, which is usually answered by someone fully familiar with the units in question. (We say 'usually' because it's hard to ignore a ringing phone!) The number is (212) 581-1424. If your question is highly technical and may require a lengthy answer, it is suggested that you call after 7PM New York City time both to save money on the call and to get the undivided attention of the answerer.

SOFTWARE:

This section begins with the assumption that the Analyzer is connected to your TRS-80 and an audio signal of nominal level is being applied to the input. If this is not the case, refer to the appendices.

IMPORTANT NOTE!

The Analyzer is designed to be controlled under LEVEL II or any version of Disk BASIC. This is because operation of the Analyzer requires the capability of machine language "calls". If you only have a LEVEL 1 machine, it will be necessary to have it modified by Radio Shack.

The Eventide Real-Time Audio Spectrum Analyzer is a combination of analog and digital circuitry designed to elicit the relative amplitudes of various frequency components of an audio signal and show them in a convenient form on the monitor CRT. The

frequency/amplitude data are also available in digital form for analysis by user-written programs. The hardware in the analyzer allows operation at very high data rates, and displaying the data in meaningful form is a very time-critical procedure. The routines necessary to do this are contained on the 1K by 8 Read Only Memory on the Analyzer board. The ROM occupies locations FC00H through FFFFH (-1024 through -1 decimal) in the TRS-80 address space. In order to avoid the necessity of memorizing many addresses, linkage to the TRS-80 operating system is by means of the USR function. By assigning values to the USR function (e.g. F=USR(1)), the machine language routines are accessed conveniently in BASIC programs. For instance, setting the Analyzer for a LINEAR display is accomplished by setting F=USR(8). For clarity in programming, we recommend a procedure whereby a variable whose name is descriptive of the function to be called is assigned at the beginning of the program, and using that variable in the USR call. Thus, a routine to perform a complete spectrum analysis would be as follows:

```
10  SCAN = 3 : GRAPH = 2
20 F=USR(SCAN)  :  F=USR (GRAPH)
30  GOTO 20
```

In several cases, the variable returned from the USR call is significant. For instance, in the case of SCAN, 'F' would contain the number of the frequency cell with the highest absolute amplitude. This is particularly valuable in voice recognition work, and saves much time over a BASIC search. In other cases, the value returned, if any, is meaningless. A summary table is provided for your convenience.

INITIALIZING

The VTU-02 contains two interdependent pieces of hardware, a 2708 EPROM and an 8255 programmable parallel interface. The hardware is configured so that the low order bit of port E5H (229 decimal) controls the appearance of the 2708 EPROM in the memory map. Although RAM may be installed at the same address as the ROM (or may not), any READ operation by the CPU will get data from the ROM when it is activated. The ON LINE LED on the Analyzer chassis shows whether the ROM is activated or not. The first step in initialization is to make sure that turning on the analyzer does not disturb the system stack. This can only be a problem if your computer has an expansion interface which contains a full 32K of RAM for a system total of 48K. If this is the case, when BASIC asks for "MEMORY SIZE?", you should reply "55000" (or lower, to accommodate any other machine language programs.)

Next, the 8255 must be instructed to turn the ROM ON. Type

```
OUT 231,144
```

This programs the 8255 control register to configure the other ports as outputs. The least significant bit (BIT 0) of the port which responds to 229 controls the ROM activation. Examples:

```
OUT 229,0 : REM TURNS ON ROM SO THAT IT MAY BE READ BY CPU
OUT 229,1 : REM TURNS OFF ROM SO THAT RAM IS ACCESSIBLE
```

NOTE: The other seven bits of this port control the GAIN of the analyzer. Care must be taken when varying the GAIN under software control to avoid disturbing the state of the least significant bit.

The final step in initializing is to tell BASIC where the ROM software begins. In LEVEL II BASIC this is done by setting the USR locations as follows:

```
POKE 16526,0 : POKE 16527, 252 : REM POINT TO BEGINNING OF ROM
```

In DISK BASIC a simpler method is to use the DEFUSR function, e.g.,

```
DEFUSR1=&HFC00 : REM USR1(N) WILL CALL FUNCTION N IN THIS CASE
```

For simplicity, we will always assume that the LEVEL II USR function (i.e., the only one available) is being used. If you are using DISK BASIC and TRSDOS, it will be necessary to substitute the appropriately numbered USR function, e.g. USR1 as in the previous example.

Following the above procedure initializes the Analyzer circuitry and links it to BASIC. All further operation is by USR calls and PEEK/POKE or OUT statements.

SCALE FUNCTION (F=USR(1))

Calling this function clears the TRS-80 screen and then fills it with the appropriate labels and scale markings. It also resets all flags (decay rate, mode, etc.) and sets the gain of the input amplifier to unity.

BARGRAPH FUNCTION (F=USR(2))

This function draws a BARGRAPH on the screen between the axes drawn by USR(1) above. There are 32 bars, each with a vertical resolution of 42 elements. The bar heights are determined by the data in locations 4200H through 421FH (decimal 16896 through 16927), left to right. Any number greater than 42 sets the bar height to be equal to 42. (Note that this routine may be used to plot ANY data set in these locations as a bar graph.)

Each bar is 1 character wide with a 1/2 character space between bars. This is important to remember if you are trying to write text under a specific bar: it will only be possible to center letters under alternate bars. The graphing process does not affect the screen outside the axes, does not affect the data in the buffer, and there is no return variable (the value of F is unchanged).

SCAN FUNCTION (F=USR(3))

This actually performs the spectrum analysis. It automatically reads the amplitude of the input signal as passed by the individual filters (and the unfiltered input as well, on the LEVL channel). See the theory of operation section for how this is done. The results of the analysis are placed in the buffer between 4200H and 421FH as described above. The various modes of operation (LOGarithmic, LINEAR, SINGLE, AVERAGE, FAST, SLOW) are active during the SCAN operation and directly affect the data placed in the buffer. The effect of each of these modes will be described under its particular USR call. The return parameter of USR(3) is an integer between 0 and 30 inclusive and is the number of the frequency cell which contained the highest signal level during the analysis. Note that there is no implication regarding the actual amplitude in this number, it is only the cell or "bin" number which had the highest amplitude. The actual amplitude may be determined by

```
10 F=USR(3)
20 CELL = PEEK (16896+ F)
30 PRINT CELL
```

The actual frequency of the cell may be determined by

```
FREQ = 20*(2^(1/3))^CELL
```

For higher cell numbers FREQ becomes somewhat of an approximation, due to the necessity of rationalizing octaves and decades, but the numbers are within the tolerance of the filters and may be used as-is if desired. If several cells have identical amplitudes and no one cell has a higher amplitude, the number returned will be that of the rightmost highest amplitude cell.

SLOW DECAY FUNCTION (F=USR(4))

When function 4 is executed, a flag is set which is accessed during the scan routine. Normally the scan routine takes the absolute amplitude of the data in each frequency cell and transfers it to the buffer location as described earlier. If the decay flag is set, the current amplitude datum in any given cell is compared with the previous amplitude of the same cell. If the NEW datum is GREATER than the OLD datum, the NEW datum REPLACES the OLD datum. If the OLD datum is GREATER than the NEW, the amplitude is permitted to decrease only by an amount set in location 41FEH (16894 decimal). When the system is initialized, the number in this location is equal to 3. It may be changed by a POKE statement as follows:

```
POKE 16894, 2: REM CUTS DECAY RATE BY FACTOR OF .33
POKE 16894, 1: REM GIVES SLOWEST POSSIBLE DECAY
POKE 16894, 0: REM GENERATES PEAK-HOLD FUNCTION, NEVER DECAYS!
```

When the SLOW DECAY function is called, the legend "DECAY" is displayed on the screen. If it is called when the value in location 16894 is 0, the legend "PEAK HOLD" is printed instead.

Note that this function operates on the SCAN data AFTER the LOG function is performed (see below). Thus the bars will appear to decay at the same rate regardless of the format of the presentation.

F=USR(4) does not change the value of F.

FAST DECAY (F=USR(5))

This function reverses the effect of function 4 above. The decay flag is reset and the value of the decay rate constant in the above POKE statements becomes irrelevant.

Function 5 does not change the value of F

AVERAGE FUNCTION (F=USR(6))

Executing function 6 places the Analyzer in the AVERAGE mode, and zeroes two areas in the memory. These areas are used to store the data developed in the AVERAGE mode. The first area is the NUMBER OF AVERAGES counter. This is a double byte located at 4220H(LSB) and 4221H(MSB) (16896 and 16897 decimal, respectively). When the AVERAGE FLAG is set by function 6, each time the SCAN function is activated, 1 is added to this double byte. Simultaneously, the data from the various frequency cells are added to corresponding double byte cells in the new buffer area. The first double byte cell of averaged data is 4222H(LS Byte) and 4223H(MS Byte). The data in these locations correspond to the sum of the amplitudes of the data in the first displayable cell (LEVL) from the time that function 6 was originally invoked. To read the data in any given double byte cell, a statement of the form

```
PRINT PEEK (16930+CELL*2)+256*(PEEK(16931+CELL*2)
```

may be used.

The data may be "normalized" by dividing the current data in each cell by the current number of averages. The value of this procedure will be made clear in the applications section. Unlike the data transferred to the standard display buffer, which are divided by 4 (so that they will fit on the display), the data stored in the averaging buffer are NOT divided. Thus, to display the averaged data, an additional division-by-4 must be performed.

Several features of the AVERAGING procedure to be noted are:

If the display is LOGarithmic, the summation is performed BEFORE the data are logged, as the "average" of logged data is mathematically meaningless.

During normal operation, the process is invisible to the user: The only visible evidence that averaging is taking place is the MODE annunciation on the screen. The overhead required by function 6 adds several milliseconds to each SCAN.

Function 6 does not change the value of F.

DISABLE AVERAGE (F=USR(7))

This function resets the average flag and thus prevents additional SCANS from changing the stored data. The value of the NUMBER OF AVERAGES bytes mentioned above is returned by the function, i.e.,

```
10 F=USR(7)
20 PRINT "NUMBER OF AVERAGES"; F : REM END AVERAGING
```

After this function is invoked, the AVERAGING legend on the CRT screen will be erased.

LINEAR DISPLAY (F=USR(8))

This function sets the DISPLAY annunciator to LINEAR, 160MV/DIV and causes the height of the bargraph created by function 2 to be precisely equal to the absolute amplitude of the data returned by the function 3 (SCAN) process. Because of physical limitations of the display area (42 elements high), the data will not be faithfully represented if the amplitude of the data exceeds this number. Circuitry limitations prevent the number from exceeding about 63, so one can regard the display as "hard clipping" about 4dB below the maximum dynamic range of the analyzer in the LINEAR mode. Please note that in almost all cases this is preferred to the obvious alternative of "normalizing" the data by multiplying by 3/4 to get it to fit within the screen height. The reason for this is that the physical characteristics of almost all signals other than pure sine waves assure that any individual frequency component is at least 4dB or more below the peak amplitude of the signal as shown in the "LEVL" cell.

The LINEAR function is executed somewhat more rapidly than the following LOG function, but the difference will usually be unnoticeable. Function 8 does not alter the value of F.

LOGARITHMIC SCALE (F=USR(9))

Calling USR(9) sets the LOG flag which tells the SCAN (USR(3)) function to compute the logarithm of the amplitude in each frequency cell and deposit this value in the buffer. As mentioned earlier, this is done AFTER any averaging is performed, and BEFORE any variable rate decay function is performed. When the LOG function is active, "LOG 1DB/DIV" will be displayed. When in the LOG mode, the screen dynamic range is 42dB, with each vertical division corresponding to 1dB.

LOG Notes: Because at very low levels even one step of A/D converter resolution corresponds to several dB, low level signals may appear "jerky" in the FAST display mode. This is pretty much unavoidable. Function 9 does not change the value of F.

CURSOR CONTROL (F=USR(10))

This function controls the cursor which moves along the bottom line of the display and which points to the bar whose frequency is displayed to the left of the screen. The LEFT and RIGHT ARROWS control the motion of the cursor. This routine must be called

continually for it to be useful. Placing it in the loop which performs the analysis is ideal. Note that the cursor does not appear until one of the arrow keys is depressed.

REMOVE CURSOR (F=USR(11))

This function re-writes the bottom line of the display, deleting the cursor and frequency legends without disturbing the rest of the display.

KEYBOARD READ (F=USR(12))

This function is provided to allow convenient interactive operation of the analyzer. Each time it is called, the return value F contains a number indicating which if any of several keys is pressed. If no key is pressed, 64 is returned. The return values are 0 through 7 for numerical keys 0 through 7; 8 for BREAK; and 9 through 12 for the ARROW keys LEFT, RIGHT, UP, and DOWN respectively.

SOFTWARE CONTROL OF GAIN

Audio signals cover an extremely wide "dynamic range", the difference in absolute amplitude between the highest level of signal likely to be present under given circumstances and that of the lowest. The analyzer can display about 42dB at any given time on its CRT. Many times, however, the signal will be so loud as to "overload" the display when it is near its maximum level. Or possibly, the level is so low that only loud portions are visible at all. The purpose of the GAIN control circuitry is to adjust the fixed gain of the analyzer so that the display gives the most usable information.

Upon initialization, a value of 128 decimal is poked into the gain control port located at E5H (229decimal). (The equivalent statement is "OUT 229,128"). This corresponds to a pre-amplifier gain of 6dB (or a voltage gain of 2). Referring to the block diagram and schematic of the unit, you will see that an input signal of nominally 5V peak to peak is required to reach full output from the LEVL detector. This gain setting would be useful for many "semi-professional" recording studios where this approximates the signal levels present on output lines. If one wanted to monitor the output of, for instance, a Hi-Fi preamplifier, a bit more gain would be desirable, perhaps 6 to 12dB. The rule is that the signal gain increases by 6dB for each division by 2 of the data word OUTput into port 229. Thus, to get an additional 6dB of gain, one would execute

```
OUT 229,64
```

An additional 6dB would be obtained by using

```
OUT 229,32
```

One convenient way to control the analyzer gain in a program is to use a BASIC variable dedicated to the gain function. The following example program allows automatic gain control dependent upon input signal level.

```
10 GAIN = 128
20 OUT 229, GAIN AND 254 : REM "AND 254" PREVENTS DISABLING ROM
30 MAX=USR(3) : REM SCAN ROUTINE, RETURN PEAK CELL
40 IF PEEK (16896+MAX) < 25 THEN GAIN = GAIN/2 : GOSUB 100
50 IF PEEK (16896+MAX) > 35 THEN GAIN = GAIN*2 : GOSUB 110
60 F=USR(2) : REM PRINT BAR GRAPH
70 GOTO 20
100 IF GAIN >255 THEN GAIN = 255: RETURN: REM PREVENT ILLEGAL QUANTITY
110 IF GAIN < 2 THEN GAIN = 2: REM PREVENT PROM DISABLE OR NO FEEDBACK
120 RETURN
```

There are a few hardware considerations of which you should be aware: OUTputting a value of zero (or any value for which the LSB=0)disables the analyzer. Jumping to locations in PROM when the prom is disabled can cause a system "crash". Using very high gain settings will affect the high frequency response of the analyzer unfavorably. If the OUT value is below about 8 (corresponding to a gain of about 30dB), a small but noticeable rolloff of response in the LEVL channel and the 20kHz filters will be noticed on sine wave tests. This becomes progressively more significant as the gain is increased (OUT value decreased). If the analyzer is used correctly, this should never be a problem, but it is something of which to be aware. The analyzer is not designed to be driven with a microphone or other low level source. A suitable matched preamplifier should be used to assure a proper signal level to the analyzer input.

SPECTRUM ANALYZER MEMORY MAP

Memory Space:

HEX	DECIMAL	USE
41FDH	16893	Current cursor position for USR(10)
41FEH	16894	Decay constant
41FFH	16895	Flag register: bit 3 - 1 = average bit 2 - 1 = DECAY bit 1 - 1 = LOG 0 = LIN
4200H - 421FH	16896 - 16927	32 bytes, one per filter, 20 Hz to 20 kHz, and LEVL
4220H - 4221H	16928 - 16929	Number of averages counter LSB, MSB
4222H - 4262H	16930 - 16994	Two byte accumulators, one per filter, 20 Hz to 20 kHz, and LEVL, used for average mode. Stored as LSB, MSB.
FC00H	-1024	Start of EPROM
FFFFH	-1	End of EPROM

I/O Space:

HEX	DECIMAL	USE
E7H	231	Control register
E5H	229	Gain and shadow control: bits 8 thru 2 - gain control bit 1 - shadow signal 0 = EPROM ON 1 = EPROM OFF
E6H	230	Filter selection, control.
E4H	228	A/D output

Note: collectively, these represent an Intel 8255A programmable peripheral interface.

FIRMWARE SUMMARY

Using the VTU02 Real-Time Spectrum Analyzer adds the following capabilities to the Radio Shack TRS-80 computer:

INITIALIZING

```
OUT 231,144
OUT 229,0
POKE 16526,0
POKE 16527,252
```

SYNTAX	RETURN VARIABLE	COMMENTS
F = USR(1)	NONE	Scale function
F = USR(2)	NONE	Bargraph function
F = USR(3)	Cell # with max. amplitude	Scan function
F = USR(4)	NONE	Slow decay function
F = USR(5)	NONE	Fast decay function
F = USR(6)	NONE	Average function
F = USR(7)	# of scans since USR(6) invoked	Disable average
F = USR(8)	NONE	Linear display
F = USR(9)	NONE	Logarithmic display
F = USR(10)	NONE	Cursor control
F = USR(11)	NONE	Remove cursor
F = USR(12)	# of key pressed	Keyboard read

See Memory Map for gain control and data buffer locations.

THEORY OF OPERATION

This spectrum analyzer is of the type classified as "REAL-TIME, CONSTANT PERCENTAGE BANDWIDTH". The applications section will give a brief introduction to other types of analyzers. In this section, we consider what the above classification means and how it is achieved in this hardware implementation.

REAL-TIME

A Real-Time analyzer is one which continuously analyzes ALL components of interest in the input signal and provides information in some usable form. This is in contrast to, for instance, the "swept filter" type of analyzer which can analyze only one frequency at a given time. There are three general methods of performing real-time analysis. The first is to do as we do here, which is to provide a large number of bandpass filters so that all components are analyzed simultaneously. The other methods involve digital capture of signal segments and their processing either by analog or digital techniques. While the latter two methods are in many cases more powerful than the large number of filters method, their disadvantages such as high cost, slow speed, or both render them unsuitable for applications in which audio spectra must be observed many times per second.

The CONSTANT PERCENTAGE BANDWIDTH consideration also enters into the choice of analysis method. The phrase means that each filter has a bandwidth which, when divided into its center frequency, yields a constant. For instance, if a filter has 3dB points 2Hz apart at 20Hz, another filter in the same analyzer should have a 3dB bandwidth of 20Hz at 200Hz center frequency, etc. This characteristic, fortunately, is one which is almost unavoidably achieved when building filters using identical hardware configurations and different resistive or capacitive values to select the center frequency. The other methods yield excellent "line spectra" but require additional processing to give constant percentage bandwidths.

A "third octave" filter set is clearly a constant percentage bandwidth type. An octave is a constant frequency ratio (2:1), and one third of an octave is also a constant ratio. Note that the theoretical method of deriving the filter center frequencies would be to start at the first frequency desired and multiply each successive filter center by 1/3 octave, or the cube root of 2. This in fact works (see formula in USR(3) description). For the convenience of us humans, it is desirable to have the filters correspond to decade ratios as well, and so the standards organizations (ISO/ANSI/IEEE and probably others) have agreed on a set of center frequencies which rationalize the disparate ratios quite nicely, and this analyzer conforms to this standard for center frequencies. The filters are identical except for the resistor and capacitor values which determine the center frequencies. They are two-pole filters with a Q of 10, which determines the bandwidth. Trim resistors adjust the Q and gain simultaneously so that small component variations can be compensated and the individual filter gains can be made almost precisely equal.

SIGNAL PATH

Although the input signal could in theory be applied immediately to the filter inputs, it is customary and convenient to provide a variable gain preamplifier. The most important reason for this is to preserve an adequate dynamic range. The filter peak-to-peak output voltage swing is limited by the power supplies to about 20 volts, and the voltage measurement capability of the unit divides the positive half of this swing into 255 even steps. If the input signal were, say, only two volts, then only 25 of these steps could be measured and the accuracy and resolution of the analyzer would be drastically limited to say the least. Variable gain in the Analyzer is implemented in a somewhat unusual fashion. A standard operational amplifier in an inverting configuration is used as the "preamp". In the feedback loop of the op-amp is a CMOS digital to analog converter whose input code is controlled by the PET. The D/A converter is of the "four quadrant multiplying" type, which means that the REFERENCE input (connected to the preamp output) is MULTIPLIED by the digital word applied to the converter. Thus, if the maximum word (255) is applied, the reference is multiplied by 1. If the input word is 0, the multiplication factor is likewise 0 and the reference input will have no effect on the output of the D/A. When the input code is 255, the DAC appears to the op-amp as a fixed resistor of about 10K ohms. As can be seen from the schematic, this gives an op-amp gain of 1. If the input code is 128, the DAC looks like a 20K resistor, and the op-amp gain is 2, or 6db. There are two reasons for the falloff in high frequency response at high gain: Note that there is a small capacitor across the feedback "resistor". This is necessary for circuit stability. As the "resistor" increases in value, the capacitor makes a proportionately larger contribution to the feedback. The other reason is that any given op-amp has a limited amount of gain at any given frequency, and at very high gains the amplifier itself begins to roll off. The gain available from the preamp is sufficient for almost all ordinary signal sources, such as audio consoles, hi-fi "aux" outputs, guitar pickups, etc. It is not sufficient for very low-level signals, such as dynamic microphones, phono cartridges, telephone pickup loops, etc. While it is

possible to build a pre-amp using the uncommitted op-amp in the analyzer, this is not recommended except in specific cases (see applications section). Close proximity to a computer is not a very felicitous environment for low-level signals, and it is recommended that an external pre-amp be used for this reason.

Careful reading of the schematic will show that there are in fact two preamps in the Analyzer, one of which is unused. The purpose of this configuration is to allow stereo operation under certain circumstances. Note that this is not supported in firmware, and serious hardware modification of ALL the filters would be necessary to provide this "feature", so feel free to use this extra op-amp in addition to the uncommitted one.

DETECTION

Once the input signal is applied to the filters, it is separated into its various frequency components. If the input contains a mixture of 100Hz and 800Hz, one would expect the two filters with those center frequencies to have outputs representative of the input amplitudes at those frequencies. One would also expect that adjacent filters (80 and 125Hz, 630 and 1K Hz) would also have output signals at somewhat lower levels, and their adjacent filters to have outputs at even lower levels. One would not be disappointed. Because these signals are AC, they may assume any value between their peak amplitudes. If we are to make a measurement of their absolute values, we must measure their "envelope" amplitude. We do this with a diode detector in a manner analogous to the diode detector in an AM radio receiver. Whenever the filter output instantaneous amplitude is more positive than the voltage on its filter capacitor, the diode conducts and charge is transferred to the capacitor. This is in effect a peak detector circuit. Note that, for the moment, there is no discharge path, and so the charge on the capacitor remains at the peak value of the AC input to the specific filter. This is a very important characteristic of the Analyzer, and we'll have more to say about this a bit later.

Also note that in one case, there is no filter between the input and the diode detector. In this case, all input signals, regardless of frequency, affect the charge on the capacitor. This charge is used to determine the height of the LEVL bar on the display.

In order to display a spectrum, we must convert the various charges on the detector capacitors to bar heights on the CRT. This is done in a number of steps:
Step 1 is scanning the various filter outputs. This is done by invoking a routine (in the machine language ROM) which sequences the PIA outputs through numbers 0 through 31. Each of these numbers is decoded, first by the 74C42 to select a multiplexer, and then by the 4 CD4051 analog multiplexer chips to select an individual capacitor. Step 2 causes the A/D converter to measure the voltage on the capacitor (0 to 255). Step 3 allows the computer to read the A/D converter output (by activating the bus transceiver chips IC's 42 and 47) and placing the data on the Apple data bus. Once the data are read by the CPU, the actual number corresponds to the voltage, which corresponds to the height of the bar. The number is stored in the buffer as described in the software section. In practice, each of these steps is performed sequentially for one filter before going on to the next. It's not quite that simple actually, but those three steps are conceptually the entire analysis.

TIMING

In point of fact, a few other actions occur between the measurements: For instance, the A/D converter is not instantaneous, and must be paced by software. If one filter were turned off simultaneously with another filter being turned on, there would be the possibility that charge would leak from one capacitor to another, etc.

For this reason, a software-controlled timing sequence is invoked for each filter. Because these routines may be performed in BASIC as well as machine language, it is very important for the user to understand the hardware requirements if he wishes to program the analyzer for exotic applications. A single analysis consists of the following procedures:
1: Inhibit the multiplexers to prevent inter-capacitor discharge. The D input of the 74C42 coming from the bit 7 of the port E6 performs this function. 2: The filter (0 through 31) is selected as described earlier, and the inhibit line lowered, allowing the capacitor to be connected to the A/D converter. 3: A very short period is allowed for the analog circuitry to "settle", and the A/D conversion is initiated by strobing the convert flip-flop from E6 (bit #6) low. 4: Another delay allows the conversion to be completed. This takes about 80 microseconds. 5: The A/D data is read as described earlier, and a resistor is connected to the multiplexer output through one section of the analog switch controlled by bit #5 of port E6. This resistor discharges the capacitor to ground, thus preparing the way for a new sample of data to charge the capacitor. This simple, unassuming step is a major feature of the unit, because it guarantees that each analysis is independent of the previous one, giving statistical validity to the output data under almost all circumstances. In effect it allows the user to control the discharge time constant (as opposed to the DISPLAY time constant described earlier) automatically by varying the rate at which the analysis is performed. 6: Finally, the multiplexer is inhibited once again and the next filter in line is selected, and the process begun again.

POWER SUPPLY

Power for the unit comes from a regulated power supply in the Analyzer chassis. The Analyzer requires +/-12V for the analog circuitry, +12 and +/-5V for the machine language ROM, and +5V for the logic circuitry. The output of a stepdown transformer operating on 115 or 230VAC is rectified and filtered to produce "raw" DC of about +/-16 volts. The + voltage is applied to the input of a monolithic 5V regulator (78M05), and to the + input of a dual regulator type NE5553. The -16V signal is applied to the - input of the NE5553. The outputs of these chips provide regulated +/-12V and +5V. The -5V required by the ROM is produced by a resistor in series with a 5V zener diode.

ADDRESS MAP

A 74LS30 decodes the Z80 address bus from the TRS-80 and asserts the NTOPK signal when these addresses are being accessed. If the analyzer has been initialized and is activated, CPU READ's to this address activate the EPROM. Another 74LS30 in conjunction with a pair of inverters decode various OUTPUT ports (see memory map) to activate the various registers of the 8255.

POTPOURRI

Two sections of a CMOS switch (CD4066) form a clock oscillator for the A/D converter, and a 74C74 synchronizes the converter with the 8255 commands. Diodes and a SIP pull-up network are used to interface the A/D (operating at 12V) to the 8255 inputs (operating at 5V).

APPLICATIONS

Let's begin by saying that this section is not, and cannot be, either exhaustive or complete. The spectrum analyzer is as versatile and protean a tool as is the oscilloscope, which it resembles in many respects. Complete books have been published about single aspects of oscilloscope usage (as they have about spectrum analyzers, for that matter), and so this section can only give a brief overview of some of the uses for which spectrum analyzers in general, and our 1/3 octave unit in particular can be employed.

TIME VS. FREQUENCY

One major concern of electronics is with the movement of electrons. Electrons move under various influences, such as the presence of electric and magnetic fields. These fields can be simple or complex, time varying or static. Our understanding of a circuit's behavior is directly related to our ability to observe, by various methods, the movement of electrons at various points. Common measuring instruments allow us to do this. The DC voltmeter allows us to quantify the relative concentration of electrons between point A and point B. The AC voltmeter does the same by giving us a statistical measurement of the same value even when the concentration is varying. A more versatile instrument, the cathode ray oscilloscope, actually allows us to visualize the time-varying behavior of the electron concentration. Most of the information which we may desire about a circuit or signal can be determined by proper usage of this device, along with our intuition or mathematics. The oscilloscope portrays a signal in graphical form, as AMPLITUDE (vertical axis) VS. TIME (horizontal axis). The spectrum analyzer portrays the identical signal in a graph whose vertical axis is still AMPLITUDE, but whose horizontal axis is now FREQUENCY. Mathematically, the two presentations are equivalent: Each can be converted to the other by an operation known as a "transform", specifically the Fourier transform. Since the two displays are "equivalent", why have (or spend money on) both instruments?

Our objective is, presumably, to learn as much about the signal under observation as possible. Let us consider the sine wave. This is the simplest time-varying signal possible. It can be completely described by two parameters, frequency and amplitude. Looking at a sine wave on an oscilloscope we see a regular repetitive curved line. Of course we recognize it as a sine wave from experience, and can easily measure its frequency by seeing how long it takes between pattern repetitions. (This time is the period, and frequency is the reciprocal of period). Of course, if we didn't recognize the sine wave, we could perform a mathematical analysis of the trace and, using the Fourier transform, come to the same conclusions to within the limits of measurement accuracy, typically 2-3% from the face of a standard CRT. This same signal applied to the input of our spectrum analyzer should produce a single vertical line (or point, depending upon the display type,) at the location on the horizontal axis corresponding to the frequency of the sine wave. The point's height is equal to the amplitude of the sine wave.

In the above case, there is little advantage to using either instrument. There really is little to know about the signal in question, and that little can be determined easily, almost by inspection. Let's consider several cases in which one or the other instrument is more effective. Assume that the sine wave is coming from the output of an amplifier which sounds bad: there is hum and distortion in the output. Recall that measurement accuracy of the oscilloscope is only about 2-3%. If the hum component or distortion component is small, on the order of 1%, we might just barely see a wavering or distortion of the oscilloscope trace, but would not be able to analyze it. On the other hand, the same signal applied to a spectrum analyzer would show spectral lines at 60Hz (or perhaps 120Hz), and at one or several harmonics of the sine wave. (We should interrupt here to point out that the 1/3 octave spectrum

analyzer is not particularly good for this type of application-this is so far only a general description of spectrum analyzer capabilities.) Chalk up one for the analyzer. Now let's consider a pulsed waveform such as that encountered in digital measurements. Assume that the signal in question is a rectangular pulse high for 1 millisecond and low for 10 milliseconds. Because we are dealing with digital circuitry, this is really what we wish to know about the signal, and the oscilloscope will give us the information by inspection. The same signal on a spectrum analyzer will give us a line at 909Hz and many more lines at increasing and decreasing amplitudes on harmonics of this frequency. In principle we can derive the amplitude and shape of the rectangular pulse from this display, but why bother? Chalk up one for the 'scope.

A major distinction among oscilloscopes and the various types of analyzers is whether they operate in REAL-TIME. This may seem a little strange to the uninitiated. For instance what is the converse? Imaginary time? No, it's "NON REAL-TIME", and the distinction, far from being some mystical or relativistic concept is: a real-time analyzer is one which analyzes all frequencies within the measurement bandwidth over the entire duration of the input signal (and, usually, presents the results of this measurement concurrently with the presence of the signal). The importance of this concept can be seen by considering a signal such as a pulsed sine wave at multiple frequencies. If our analyzer is measuring frequencies one at a time, and a pulse comes in on a different frequency, it will be ignored. Thus the display may or may not contain relevant or complete information. The main application for non-real time analyzers is the examination of continuous signals, such as the example above of a sine wave which had superimposed hum and distortion. Since such analyzers typically have one filter, its characteristics are more readily changed to conform to a specific measurement requirement than would be possible in a multiple filter unit. Audio program material is clearly time varying in frequency and amplitude components, and so is best handled by the real-time analyzer.

One more note on "real-time": Most oscilloscopes are nominally real-time devices, in that in principle the entire signal is observable. In point of fact, this characteristic is almost totally unusable on a time-varying signal. The volume of raw data to be analyzed and integrated is hopelessly immense. For instance, storing one second of data in the form it is normally shown on a 'scope would require (assuming an audio signal) about 50,000 bytes of data. This same signal, if it were displayed on a 1/3 octave analyzer with 15 updates per second, would require only 480 bytes, a compression ratio of over 100 to 1. This is not to imply that the analyzer may necessarily be used as some fantastic data compressor: information is lost in the process and in this case may not be entirely recovered by mathematical manipulation. The beauty of this is that the lost information is of little consequence and is better off gone anyway. Without becoming too philosophical, let's just say that science, intelligence, mind, or whatever you want to call it, is the ability to bring order out of chaos. The third-octave, real-time, audio spectrum analyzer does an excellent partial job of this. When installed in an intelligent controller such as the PET, Apple or TRS-80, a further reduction of data may be performed using the computer's mathematical capabilities. In fact, as one application will show, about 24 hours of data from multiple input sources can be compressed into a few thousand bytes of data. Needless to say, data in any quantity is useless unless somebody needs it and can interpret it. The following pages can be used as guidelines for possible uses of your Analyzer. If you have any unique applications not mentioned herein, we would be pleased to know of them and, with your permission, write them up in future addenda to this manual.

GROUND RULES

The applications listed here are suggestive. We are not supplying complete, proven programs for one simple reason: It is not economically feasible. Writing a program to do any of the things mentioned below could take from minutes to days. TESTING that same program could take from hours to months! As the user community for this product expands, we expect that program exchanges and translations will take place. We are currently supplying RTA's for the PET, the TRS-80, and for the Apple, and may do so for other machines as demand develops. Again, economic considerations are important. If the products sell in enough volume to warrant it, we may well have someone working full-time on user communications. If they don't, we'll still be here to answer your letters and 'phone calls!

APPLICATION #1: Watching and waiting

This is the simplest yet most important application of the Analyzer. Load the interactive program into your TRS-80 and connect the audio input to the radio. Adjust the gain to get a full scale display, and sit there and watch it for a few hours. If you have an oscilloscope, connect it in parallel with the input and watch BOTH displays. This, as will nothing else, give you an appreciation of the capabilities and limitations of the analyzer, as well as a very basic understanding of the components of an audio signal.

Frequency ranges of human voice and some instruments

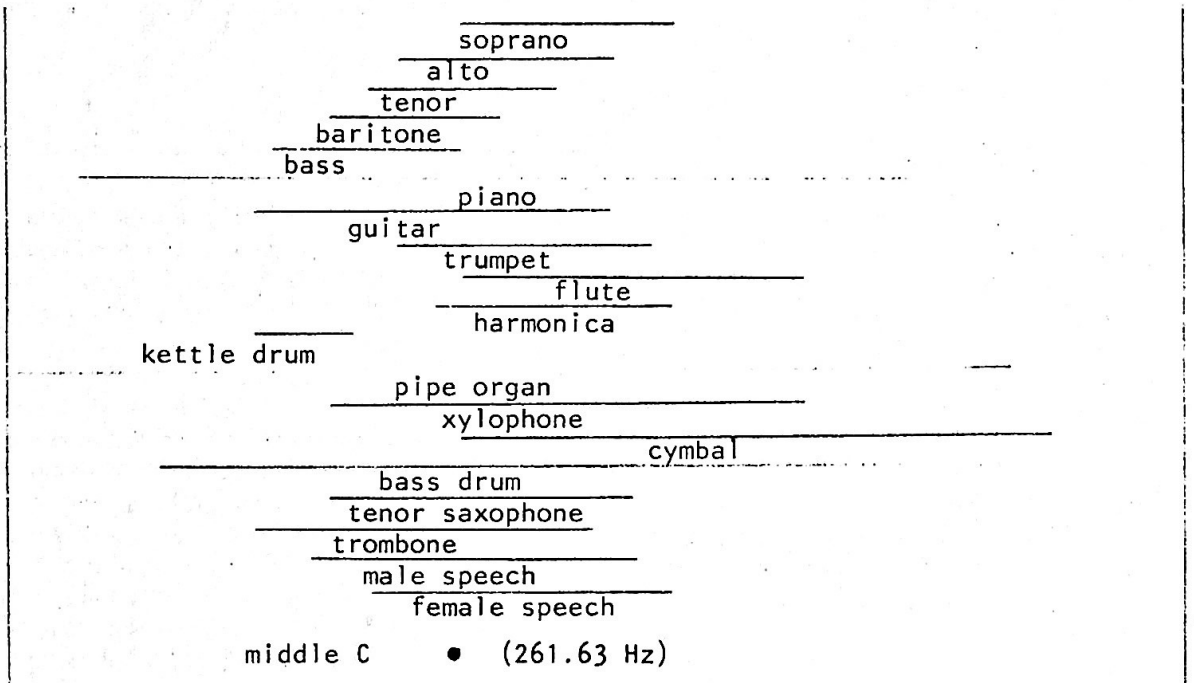

20 Hz	25 Hz	31 Hz	40 Hz	50 Hz	63 Hz	80 Hz	100 Hz	125 Hz	160 Hz	200 Hz	250 Hz	315 Hz	400 Hz	500 Hz	630 Hz	800 Hz	1000 Hz	1250 Hz	1600 Hz	2000 Hz	2500 Hz	3150 Hz	4000 Hz	5000 Hz	6300 Hz	8000 Hz	10 kHz	12.5 kHz	16 kHz	20 kHz

APPLICATION #2: Recognition

This is a very broad category, but a few hints will imply more ideas than can be implemented in a year of experimentation. Any given sound, noise, speech utterance, etc. has what is known as a "spectral signature". This is not ghostly handwriting, but rather a description of the spectrum analysis of the signal over some convenient time period. An explosion can be differentiated from a policeman's whistle by noting the predominance of low frequencies in the former. A policeman's whistle can be differentiated from a human's whistling by noting that the former is typically fixed in frequency while the latter is step or continuously variable over short periods and one to three octaves. The human ear is exquisitely sensitive to small differences and nuances in spectral signature. It is one of the most critical and interesting problems in the Artificial Intelligence and Speech Recognition Communities to determine how this is possible and to duplicate and surpass human performance using the far more precise electronic instruments available. Much progress has been made so far using ultra-sophisticated computers and fast Fourier analyzers, and this type of performance cannot be expected from our RTA system. Nonetheless, the spectral signatures of many common sounds, including speech and music, are susceptible to analysis. Depending upon the user's requirements, one might write programs to analyze and differentiate between any or all of the following:

EMERGENCIES:
Many emergency conditions are frequently characterized by certain emitted sounds: a baby's crying, the roar of an approaching tornado, the shriek(?) of a smoke detector, the scream of a damsel in distress, the burbling of someone drowning in the pool, the ding ding ding of an approaching ice cream truck (for some of us this IS an emergency, albeit a pleasant one). You get the idea. Any or all of these sounds should be distinguishable with creative BASIC programming and the RTA.
CONTROL:
Voice control of appliances by the bedridden or the lazy should be possible as well. Although we doubt that a large vocabulary can be recognized with great facility and accuracy, it should be possible to program frequently used control words (stop, start, left, right, on, off) plus a few specialized ones (higher, lower, thank you) and have them quickly and accurately recognized. Remember that the normal capabilities of the computer are untouched by the Analyzer-the other interface slots may be wired as always to devices to be controlled.
RECOGNITION
An obvious subset of this function is to accept utterances from only specified talkers. Again, for a small number of words it should be possible to achieve reasonable accuracy at this function.
Voice is not the only thing to be recognized. We feel that there would be immense advantage in having a device which would recognize disco music and, for instance, be able to throw a relay! Which way? Up to you!

The above by no means exhausts the recognition possibilities. Here's a simple one (credit to Robert A. Heinlein), how about a door opener that responds to a whistled "Yankee Doodle"? Just don't come home with your mouth full of crackers.

APPLICATION #3: LEVEL CONTROL/NOISE WEIGHTING

The ear responds differently to different frequency sounds. Equal amplitude signals can vary from almost inaudible to piercing. This is among the reasons why certain program selections will sound objectionably loud compared to their surroundings even though the broadcaster swears that his meters read identically. This phenomenon has lead to the acoustician's practice of "weighting" program and noise-like signals so that each frequency's contribution is roughly equal to its location on the equal audible amplitude contour, rather than being added in with all other frequencies in a

standard root-mean-square addition. If one desires such a measurement, the flat response of the RTA may be modified to follow ANY arbitrary curve, be it RIAA, "C", NAB, "A", CCIR, or a profile of the New York skyline. This is done by preparing a data table (using DATA statements) which is applied to the data stored in the buffer after the USR(3) SCAN but before the USR(2) bargraph. The procedure is similar to the following:

```
10 DATA W1,W2,W3,...,W31 : REM CURVE WEIGHTS
20 FOR I=0 TO 31 : READ WT : OLD=PEEK(16896+I)
30 POKE 16896+I, OLD*WT : REM FOR LOG SCALE, POKE 16896+I, OLD-WT
40 NEXT I
```

The weighting function must be consistent for LIN or LOG: if you're using a linear display, the Wn should be a fraction (0 to 1), if using a logarithmic display, the Wn is the difference in dB between reference level and the actual level of the curve at the point.

Once you've modified the measured signal by the appropriate loudness contour curve, you might desire to use the information to control the signal source, probably to keep the loudness level constant, or possibly to compensate for changes in external conditions such as the size of the listening audience. A simple way to do this is to connect a multiplying D/A converter to the user port (see the Analyzer schematic to see how to hook up the D/A). In effect this gives you a programmable gain channel identical to the Analyzer variable gain pre-amp. If you use a decent op-amp, such as an LM301 or one of the new Bi-Fet units, extremely good performance should be obtained. Multiple DAC's may be connected in parallel digitally to control stereo or quad sources.

Even further control may be obtained by using the machine language clock in Disk Basic or a separate hardware clock. Yet another weighting function may be programmed, one which will give a loudness contour over time (and date). For the entrepreneurially inclined amongst you, we recommend approaching airports, industrial plants, and public places afflicted with Musak (Trademark of Musak, of course) with a complete system which will allow this type of versatile sonic control. The flash factor alone should be worth a few kilobucks!

Application #4: PEAK HISTORY

A frequent requirement among audio professionals is to determine whether any preset peak amplitudes have been exceeded. For instance, when one is manufacturing a master recording lacquer, a mechanical process in which a cutting stylus physically removes material from a surface in a spiral pattern is employed. A compromise between the spacing of the spiral (closer spacing gives longer playing time) and the maximum signal amplitude (larger amplitude allows a better signal to noise ratio) is necessary, and if the compromise is momentarily infelicitous, one can have the stylus cut through from one groove to the next. The physical peak-to-peak amplitude is a function of frequency, so monitoring the absolute electrical amplitude of the input signal does not necessarily imply whether a crossover has occurred.

Ideally, one would like to know the peak amplitude of the input in various frequency bands, and the Analyzer is able to present this information. All one need do is connect the analyzer to the audio line feeding the record cutting equipment and set it in the PEAK HOLD mode by using the SLOW DECAY mode (USR(4)) after a POKE (16894,0) statement, which sets the decay rate at zero. Now, the result of any analysis which results in any cell having an amplitude greater than its amplitude in any previous analysis will have the new data substituted for the old in that cell only. If the

analysis is started at the beginning of a mastering operation, the screen will, at the end of that operation, give an immediate indication of whether a crossover is likely to have occurred, and what steps to take to prevent it. For instance, if the signal amplitude in the 80Hz band is disproportionately high, one could employ a standard third-octave graphic equalizer to reduce it. If the amplitude in many bands is too high, a general reduction in level is indicated. Naturally, the Analyzer does not know what frequency/amplitude combination will result in crossover in any given cutting system, and this must be determined before the information presented by the Analyzer can be utilized. Once this is determined, one can write a program which could relate playing time, amplitude, frequency, type of music, etc. to crossover probability and have the computer actually pinpoint in time when (how far into the disk) it may have occurred so that microscopic examination can verify the problem, or perhaps prevent wasting a master.

The above is a very practical application for the Analyzer. The peak-hold function may be used to store maxima relating to other phenomena. Physical events of unpredictable occurrence lend themselves to this technique. Testing to destruction, maximum noise level over time in manufacturing plants and along roadways, investigation of psychic phenomena (well, not necessarily physical!), determining level requirements for emergency communications systems, and locating intermittent faults in electronic equipment are among those which come to mind. Remember that the Analyzer can detect when a signal exceeds a certain level and generate an electrical signal to summon the operator or request other attention.

APPLICATION #5: SNEAKY RADIO

If you're a radio station, tune in. If you're not, you might find this an interesting sidelight into the "real world". All a radio station has to sell is the presence of its listeners, and theoretically the only way a radio station can get those listeners is by keeping them interested in the programming and the sound of the station. Because of the high stakes in this game, an incredible amount of talent and money is invested in making a station sound as "good" as possible, good being defined primarily as "loud" so that someone casually tuning across the dial will instantly stop, and to a somewhat lesser extent, as distortion free as the previous constraint will allow, so "listener fatigue" doesn't set in too soon. To this end, all manner of processing equipment (including some manufactured by Eventide) is employed to compress, expand, flatten, broaden, contour, adjust and squish the sound of records, dj's, advertisements, and even silence into the mold the program director (or occasionally the chief engineer) desires.

When this works, as it frequently does, a certain station may increase in the "ratings", causing competing stations in the same area to: 1; wonder what it is that their competitor is up to and 2; wonder what they can do to duplicate it and get their listeners back. This is where the Analyzer comes in.

There are two main differentiae of broadcast signals (apart from program content, of course). They are frequency content and amplitude content. We have already utilized the Analyzer to graph frequency content. We can also, with a minor hardware modification, generate information on amplitude content. The mathematical form of the information we wish to determine is the "probability density histogram", a quantized form of a probability density function. The hardware modification consists of removing the peak-hold capacitor on the LEVL channel so that sampling may be performed more rapidly.

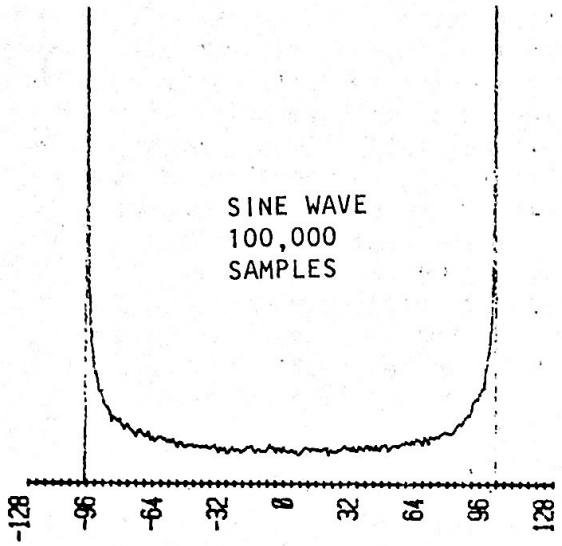

PROBABILITY DENSITY FUNCTIONS OF SOME COMMON SIGNALS

The vertical axis is the number of occurrences (not to scale).

The horizontal axis is the BIN #, corresponding to the amplitude in arbitrary units, but scaled to the range of the A/D converter.

As can be seen, a sine wave spends much of its time near the amplitude extremes. Sawtooth, triangle, and white noise signals have an equal probability of being anywhere.

The three sine wave graphs show the necessity of having a statistically adequate number of samples.

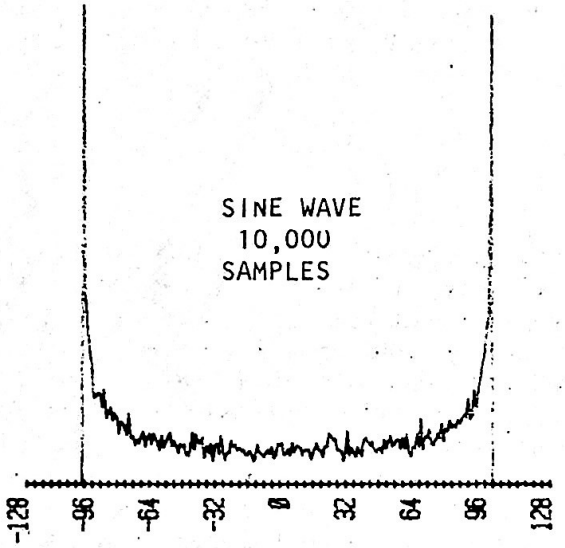

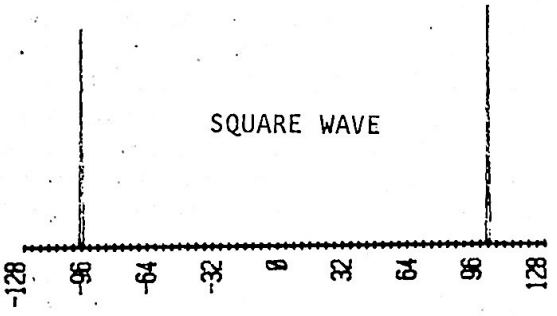

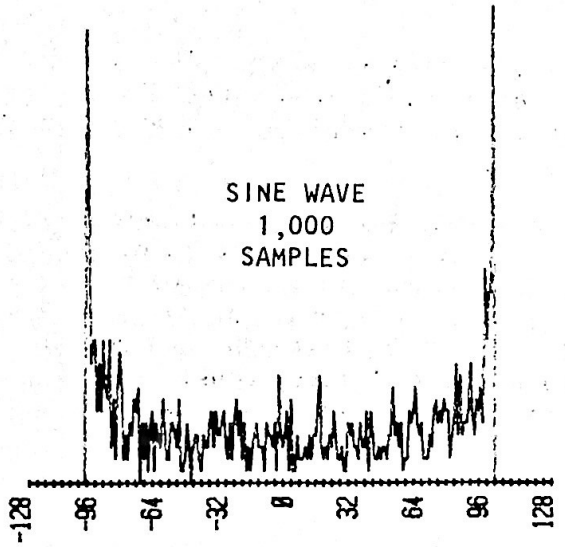

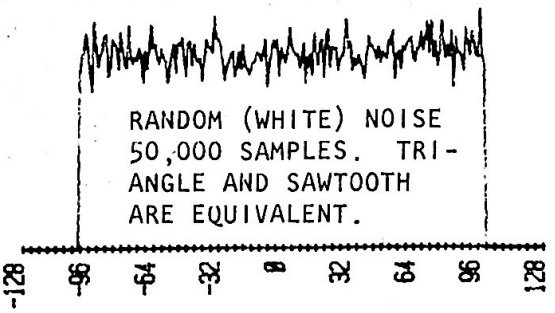

Assume one station is broadcasting a signal that another station wishes to emulate. One can monitor the putative emulatee and analyze the signal during a fixed interval when that station is playing a specific record, advertisement, or some other pre-recorded signal that is accessible to the would-be emulator. The same signal can then be analyzed as recorded, and the differences between the analyses gives some very good clues as to what processing the station is employing. In performing the analysis, two techniques will be helpful. The first is the AVERAGING mode built into the analyzer. Because we are interested in frequency content over a long period, we should take the various spectra and add them together and then divide by the number of averages to form a composite of the frequency response curve. Any individual analysis is conditioned by the idiosyncracies of the input signal at a given instant, while a group average is indicative of the response of the system as a whole, as well as the average frequency content of the program source. Because this latter characteristic is a constant for both stations, subtracting the averages point-by-point gives the equalization curve of the station in question.

The amplitude histogram gaves a general idea of the "loudness" of the station in question. By looking at a signal such as a sine wave on an oscilloscope, one can get an intuitive feeling for the probability density of the signal. Divide the vertical axis into 256 cells and examine the amount of time the CRT dot spends in each cell. If the signal is centered and less than full scale, you can see that, above and below the peaks of the sine wave, no time at all is spent. We would assign a numerical zero to all the cells above and below the signal. Again looking at the signal, we can see intuitively that little time is spent in the center cells since the dot is moving through them very rapidly. These cells would be assigned a number greater than zero, but much lower than 1. As we get closer to the peaks, we see that the signal spends relatively more time in the cells in question. Thus a graph of the probability density function of the sine wave would look like a teacup with walls asymptotic to the vertical. Graphs for other common signals look like straight lines (triangle and sawtooth), and two vertical lines equidistant from center (perfect square wave). The probability density function of both program and noise-like signals is typically a Gaussian "bell shaped curve". Assuming that one wants maximum "loudness", a reasonable assumption would be that a curve biased towards the outer extremes, indicating longer periods at higher amplitudes, would be more likely to fill the bill than one with a much more prominent central peak. Of course, the limiting case of this is a 100% "clipped" signal, where everything comes out as a square wave. While this is quite loud, it sounds horrible and only the most indiscriminating listener would prefer this signal to a relatively undistorted one. Relating again to the analysis, the extent to which the histogram taken from the received signal differs from that of the reference signal gives an indication of the processing being performed by the station in question.

Here are a few hints on actually programming the analysis: The larger the number of averages and amplitude samples, the greater the statistical validity of the procedure. It should be easy to get enough averages during the typical 3 minute song or 60 second commercial. The amplitude histogram might be better performed in machine language coding although a sufficient number of samples could probably be obtained by interrogating the A/D converter in BASIC while looking at the modified LEVL channel. Machine language is ideal for this type of analysis because the A/D requires about 80 microseconds per conversion and may be operated almost continuously since the routine to read and store the data (using STA (LOC,Y) for the 6502 or LD (HL),A for the Z80) operates in a few microseconds at most. Theoretically, using the A/D converter without a "sample and hold" stage will give incorrect results, especially with high frequency input signals. As a practical matter, this is not too serious for this type of analysis although it would be disastrous if one were to try to digitize directly the signal for later reproduction.

APPLICATION #6: MULTICHANNEL LEVEL INDICATION:

While we're performing simple hardware modifications, here's one for the recording studio, home recordist, process control engineer, transmitter watcher, or anybody who has to monitor a lot of meters at once. Grab a handful of ordinary silicon diodes (1N914 is the most common but anything will do) and connect them so that their cathodes are connected to the cathodes of the rectifiers already on the board. Do this for each channel of DC voltage you wish to monitor. These diodes essentially form analog non-exclusive OR circuits with the diodes already present, and will be multiplexed and scanned just as are the outputs of the filters. (If you use this scheme, be sure to apply no input to the Analyzer at the same time). As with all hardware modifications, be sure to observe the limitations of the circuitry. The output of the diode is periodically loaded with 4700ohms, which may introduce transients into the driving circuit, and permanently loaded by .01 microfarad filter capacitors. The measurement level is fixed at about 0 to +12Volts. The circuit must be protected from any input exceeding +12. If your driving circuit can exceed this voltage it will be necessary to connect a series resistor and a protection diode (from the drive, anode, to the +12 supply, cathode), so that any excess input current will be diverted to the power supply rail. This protection is absolutely necessary, by the way. Exceeding the rated power supply voltage will destroy the multiplexing IC's. The inputs are relatively insensitive to static charges because of the filter capacitors after each diode.

Buffered audio signals may be connected directly to these new "inputs", as may any DC signal within the input range. The resolution of the A/D converter is about 40 millivolts, so signals smaller than 10V may be measured with some loss in resolution. No specific accuracy is guaranteed, but it will be typically very good, in that the power supply is regulated and the A/D takes its references from ground and +12V.

Using external inputs as suggested above will require software changes as well. For instance, it is unlikely that the built in frequency scaling will be meaningful. Likewise, the amplitude range may be restricted. The graph axes drawn by USR(1) may be reproduced or modified in BASIC, albeit at a much lower speed, and changes may be made in the BASIC routines with impunity. Finally, if you plan to make this addition permanent, it is recommended that good wiring practices be followed. Add a connector in some convenient location and wire the anodes of the new diodes to it using the shortest path possible. Remember that any lead length between the input of the diode and the connector will act as an antenna and pick up any high frequency radiation present, giving a small, variable, and annoying DC baseline.

APPLICATION #7: DATA STORAGE

Unlike any spectrum analyzer for less than 10 times the price, the Eventide RTA cum computer has the capability of electronically storing measurement data on the program storage cassette.

This facility will permit the user to collect and tabulate data for a long period of time. Probably the most efficient method of storing the data, at least in terms of tape usage, is to define a string variable with a fixed format, such as 32 bytes for a title, 64 bytes for the data (stored as two byte hex values), and a few extra bytes for measurement data such as time, gain settings, date, and possibly a program-generated checksum. The individual data points can be stored in the string by PEEKing the buffer locations (see memory map), converting to hex, and putting the hex values at the end of the string.

Use of this feature will enable one to build up a library of spectra of various types, such as hall responses (for touring rock and roll bands), identical words from different talkers (the speech researcher), frequency responses of various radio

stations (the broadcasting consultant), etc. Of course data loaded in from these tapes may be compared with data being currently obtained to check that things haven't (or have) changed in the interim.

APPLICATION #8: Room acoustics

Almost last, and not at all least, is the primary application for the 1/3 octave analyzer, room tuning. Ideally, music should be listened to in an environment which does not significantly change or "color" its characteristics. Colorations include the addition of reverberation and the modification of frequency response. Although modern hi-fi tuners and amplifiers are capable of almost perfect reproduction, this cannot necessarily be said of phono cartridges and definitely cannot be said of speakers. Far worse is the listening environment, where imperfect speakers are coupled to completely uncontrolled acoustic spaces (rooms). For this reason, a correction factor is frequently added in professional installations. That correction is inserted by means of an analogue of the RTA, the third-octave EQUALIZER. These devices have, essentially, one knob corresponding to each bar on the Analyzer. When they are set correctly, the knobs will probably be about as high above a given center line on the equalizer as the bars are below an arbitrary center line on the analyzer. In other words, the equalizer compensates for inaccuracy in the equipment or environment by bringing up frequencies that are attenuated and attenuating those that are too high. While most studios and many hi-fi buffs have equalizers in their systems, the weak link has always been the analyzer, a very expensive (till now), and infrequently used (we hope only till now) device, which was unaffordable by most individuals and organizations.

Now that the analyzer problem is solved, let's look at the method. The simplest and most obvious method (which, of course won't work), is to put a sine wave at a fixed level into the equipment, and measure it somewhere in the room. It won't work because standing waves are created within the measurement space which will make the measurement extremely sensitive to the position of the microphone within the room (and inconsistently so: the position will have to change with frequency changes). The solution is to use a form of test signal called "pink noise". White noise is of a type which contains equal components (averaged over time) of every frequency. Because there are "more" frequencies in higher octaves, this form of signal has an additional 3db of power in successively higher octaves. Because we desire equal power within each filter's bandwidth, we connect a "pink noise filter" (see Appendix #3) which reduces the amplitude of the white noise by the requisite 3db per octave throughout the audible frequency range. White noise is readily generated by either analog sources (such as reverse-biased zener diode or base-emitter junctions), or digital sources, such as "maximal-length, pseudo random sequence generators", which are a bunch of shift registers with feedback from several outputs to the input via exclusive-OR gates. An even simpler realization of this circuit is the National Semiconductor MM5837 white noise generator chip (see appendix again). The circuit shown will furnish a perfect test signal for room alignment, so that all that is needed is a proper microphone and flat pre-amp. Of course a calibrated condenser microphone should be used for professional applications, but some of the newer electret units available at reasonable cost are almost as effective.

When making the measurement, you will probably find it effective to use the averaging mode of the Analyzer. The test signal is noise, which by definition has no fixed amplitude. Therefore, although the average amplitude in any frequency band will be constant, individual readings will vary by several dB. This effect is especially prominent at low frequencies. Displaying the result after several averages will reduce the uncertainty by a large amount. (Theoretically, N averages will reduce the uncertainty by the square root of N, so that 64 averages will improve the measurement by a factor of 8. In order for this to hold true, the averages must be statistically independent. This dictates an averaging period of at least twenty seconds, so that

the 20Hz filter will have at least 200 milliseconds between measurements. As the filter frequency increases, the damping time decreases proportionately so that in most analyses at normal speeds this factor need not be considered.) Perfectly usable results are obtained without averaging, but the display may be a bit harder to interpret at the low-frequency end.

One suggestion: As a consequence of Murphy's law, it will never be possible to get a completely flat room curve. If you find that you must go to extremes with the third-octave equalizer to achieve a "flat" response, there may be other problems.

Experiment with speaker placement and absorptive materials such as drapes, or in a studio environment, gobos and bass traps, or even different speakers. It has been the writer's personal experience that excessive "corrective" equalization is far worse than a less-than-ideal room curve. There is a lot more to say on this subject, and no more space. We would like to hear from those of you who try this procedure under less-than-ideal conditions, such as home environments and, especially, using "non-professional" microphones. We will try to pass on that information which we feel might be useful.

APPLICATION #9: MEASURING ROOM REVERB CHARACTERISTICS

Another room characteristic which affects the quality of a listening space is its reverberation time. This is normally measured by exciting the space with a noise-like signal and measuring the time it takes for the detected level of the signal to go from threshold A to threshold B, 60dB lower. This measurement, called the "RT60" time, is typically .1 to 10 seconds, depending upon the size and the composition of the walls, ceiling, floor and other reflective surfaces (seats, balconies, etc.) "Natural" reverberation is characterized by an exponential decay of the sound level (constant number of dB per second), and a relatively faster decay of higher frequencies. Anyone who has ever sung in a bathroom has heard "unnatural" reverb, and we know how unpleasant it would sound for musical listening. Analysis of the bathroom environment would undoubtedly indicate the need for sound absorbing material on the walls to eliminate the tremendous mid-frequency resonances. Some other techniques, frequently impractical in the home, include modifying surface parallelism (such as building rooms with no two surfaces parallel, and adding layers of acoustical tiling in the oddest places.

Information on what needs to be done (although not how to do it) may obtained fairly easily from the Analyzer. Step 1 is to excite the room with a burst of noise from the pink noise generator described earlier. Then using a microphone and preamp, repetitively measure (start with about 100 millisecond intervals) the amplitude in the LEVL cell, and place the data in the buffer so that a bar graph may be drawn. The bar graph will give a picture of the decay characteristics of the room at all frequencies. Adjust the sampling interval as indicated, longer for larger rooms, shorter for smaller rooms. Note that although the analyzer instantaneous dynamic range is less than the 60dB required by the standard measurement, one can either switch gain halfway through the measurement, or extrapolate the decay (assuming that it is exponential) over 60dB, most easily by doubling the time interval required for a decay of 30dB. To determine the decay time in any frequency range, simply use the data from a particular filter or filter group.

We conclude the Applications section with #10, a reprint of the article "SOUND INSULATION REQUIREMENTS FOR RECORDING STUDIOS", by Michael Rettinger, which appeared in April 1979 "Recording Engineer/Producer. This article should give the entrepreneurs reading this still more ideas.

Sound Insulation Requirements for Recording Studios

by **Michael Rettinger**
Consultant in Acoustics
Encino, California

In planning the construction of a new recording studio, or even in evaluating its acoustic climate after completion, a 24-hour noise exposure level survey is desirable. This may reveal unexpected rumbles during a certain hour of the day, periodic clinks and clonks of machinery near the room, and other equally objectionable acoustic disturbances.

Unfortunately, the making of such a prolonged sound recording requires over 20 miles of magnetic tape running at 15 ips, plus personnel for the three eight-hour shifts. Additionally, to secure statistical information on the project, the recordings have to be either re-recorded onto a graphic level recorder to make the temporal noise level variations visible, or else they have to be evaluated in some other time-consuming manner.

To economize the process of noise exposure level surveys, micro-sampling has come into vogue. Since very rarely the important noise peaks have a predominant frequency above 6,000 Hertz, so that high-fidelity recording is not necessary, a tape speed of 1-7/8 ips can be employed in the survey, to save both tape and to reduce the number of reel changes per hour. Also, the magnetic tape recorder does not have to operate continuously, but for only a given percentage of time. Thus, when the recorder is "on" only 10% of the time, say 1 hour and 12 minutes in 12 hours, the noise history may be recorded on only 675 feet of tape, which, when 0.001" thick, can be put on a standard 5" diameter reel. Similarly, a "one-hour" cassette operated 10% of the time can accommodate 10 hours of noise recording.

So as not to miss exceptional noise level changes in an hour's period, it is best, in the 10% recording technique, to record for six seconds continuously every minute, instead of recording continuously for six minutes every hour.

The switching device, known as intervalometer, must not deactivate the electronics of the recorder, since this would introduce a significant time lag in the six-second intervals of recording, but must apply a brake unit to the tape transport mechanism which is able to start and stop the medium instantly.

The accuracy of the micro-sampled results are remarkably good. The micro-sampled decile noise level exceeded 50% of the test period, or L_{50}, is less than 0.5 decibels different from the L_{50} obtained by continuous recording, when the noise is chiefly street traffic disturbances. The accuracy would not be as good if many short-term high-level signals had to be evaluated, such as foghorn whistles, gunshots, siren blasts, thunder, etc.

In practice, when the tests are to be carried out in the open, the small reel-to-reel or cassette recorder and associated intervalometer may in the evening be hidden in a bush on the site, or if overnight parking is permitted there, the equipment may be locked in a car, with the extension microphone hanging unobtrusively on the outside. Transportation of the investigator to and from the site may be by an associate or by taxi. When the site is large, this investigator has deployed several tape recorders about the lot to obtain the temporal noise level variations at several locations.

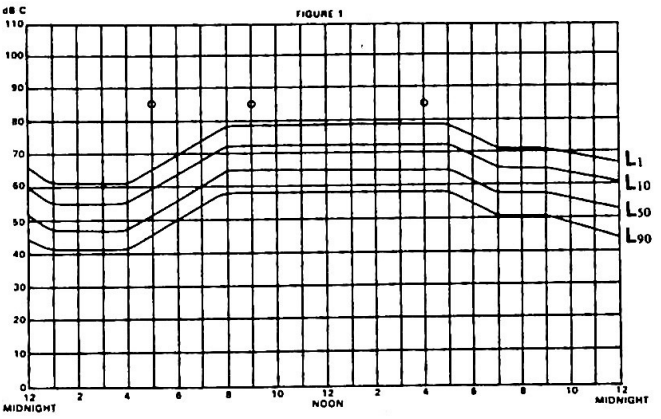

Figure 1 shows a 24-hour noise exposure level survey made at a proposed studio site. The designations L_1, L_{10}, L_{50}, and L_{90} stand for the C-weighted decile noise levels exceeded respectively 1, 10, 50 and 90 per cent of the time. In environmental assessment work the A-weighted sound levels are often used in such surveys, to gain a measure of the annoyance which the acoustic disturbances may produce on the dwellers of the area. This is of less importance here, since in studio work one is more interested in the amount of sound insulation required to achieve a specified interior noise level limit from a knowledge of the exterior sound pressure level spectrum.

When a statistical distribution analyzer is not available for determining the decile noise levels, they may be obtained by adding all the time intervals between two specified noise levels, as obtained from the temporal noise level variation chart, as illustrated on Figure 2.

A noise histogram shows the precentage of time that the noise level at a given area is found between two set limits. This amplitude distribution is by means of rectangles whose widths are the noise level ranges and whose heights are the temporal percentages.

The question arises now what can be done towards calculating the required sound insulation of the exterior boundaries of the planned studio after the noise exposure level survey has been conducted at the prospective site. First, from the magnetic tape used to record the temporal noise level variation, the spectrum of the prevailing noise should be obtained. This may be done by connecting the output of the

FIGURE 2

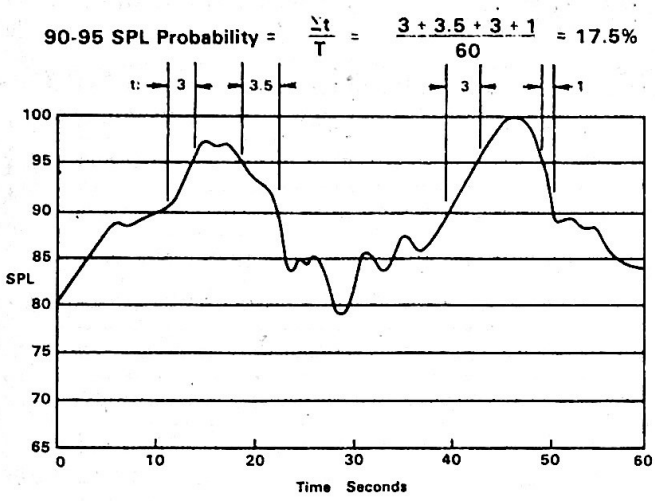

recorder, through a variable bandpass filter, to a graphic level recorder or other instrument to obtain the temporal sound pressure level variations in the various octaves or third-octaves contained in the variable bandpass filter. We obtain thus a series of charts, much like that of Figure 1, each chart pertaining to a different frequency band. Again we must perform a statistical evaluation of the octave components, that is, decile noise levels, L_1', L_{10}', etc., where the L primes pertain to the various octave sound pressure levels.

From a practical point of view, it is not necessary to consider $L_{.1}$ levels, that is, decile noise levels exceeded by .1% of the test period, because such a small percentage refers to only the noise levels in 3.6 seconds within an hour. When the standard deviation of the noise (a measure of the spread about a mean value) is 6 dB, the noise level difference between $L_{.1}$ and L_1 is 4.5 dB when the normal distribution of the noise levels is near Gaussian. To use $L_{.1}$ as the criterion value for the exterior noise instead of L_1 would require twice the surface density of the walls to achieve a desired interior noise level limit. This means twice as thick walls, since surface density refers to the mass per unit wall area.

Figure 3 shows the series of curves known as PNC graphs. Each curve is a code for specifying the noise level characteristic, or spectrum, of the noise in a room. In sound recording studios PNC-20 is generally desired, which is equivalent to an A-weighted sound level of 30 dB-A.

Figure 4 represents a graphic solution to the problem of how to determine the required sound insulation of the exterior walls of a planned studio. The solid line is the prevailing noise level spectrum during 1% of the time, the so-called L_1', as discussed previously, and is labelled Curve A. The dashed line is the PNC-20 characteristic desired for the interior of the studio, and is described as Curve B. Curve C is the difference between Curve A and Curve B, and constitutes the required sound insulation characteristic of the exterior boundaries of the planned building. □□□

Recommended Background Noise Criteria for Various Rooms

Type of Enclosure	PNC	A - Weighted Sound Level
Recording Studios, Concert Halls	15 - 20	25 - 30
Theatres	20 - 25	30 - 35
Homes	25 - 30	35 - 40
Offices	30 - 35	40 - 45
Stores	35 - 40	45 - 50

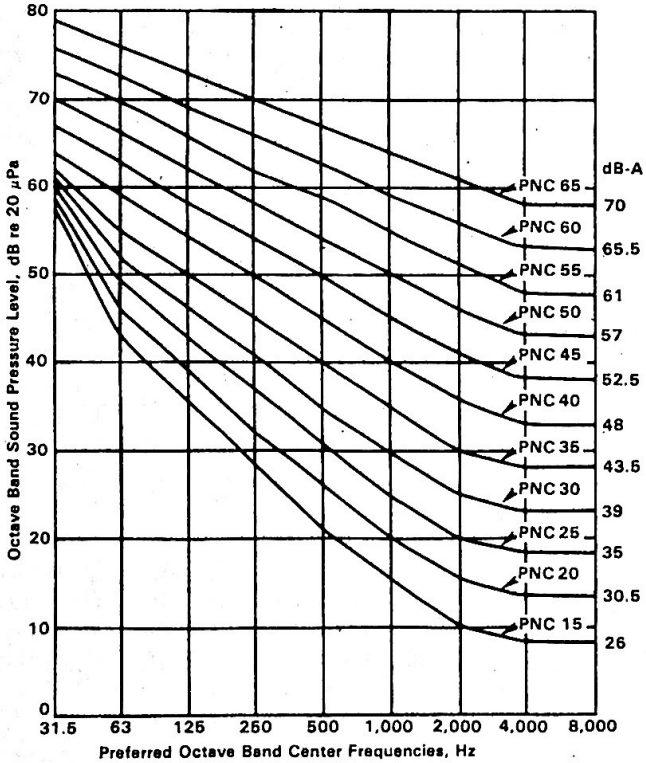

FIGURE 3: PREFERRED NOISE CRITERIA CURVES

Each curve is a code for specifying the noise level characteristic, or spectrum, of noise in rooms.

FIGURE 4

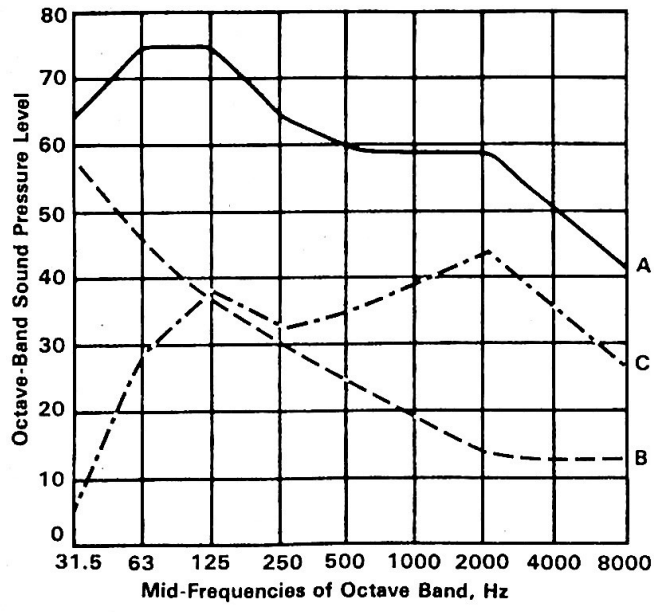

Curve A: Decile noise level spectrum existing 1% of time.
Curve B: Preferred noise criterion PNC-20, equal to 30 dB-A.
Curve C: Required sound insulation characteristic of studio (curve A minus curve B.)

APPENDIX #1: INSTALLATION INSTRUCTIONS

The VTU-02 spectrum analyzer connects to the expansion port at the rear of the TRS-80 computer chassis. It has its own power connection and does not use power from the computer or any computer accessories. As with all TRS-80 accessories, it is highly recommended that a noise-free power source be employed.

The cable coming from the analyzer plugs into the expansion port without twisting (providing that the analyzer is print-side up!) Any accessories presently attached to this expansion port may then be plugged into the equivalent port at the rear of the analyzer. Audio input to the analyzer is via the 1/4 inch phone jack mounted on the analyzer chassis. AC power should be applied to the analyzer either before or simultaneously with power to the computer. Two LED indicators are provided, one to show that AC power is applied, and one to indicate when the analyzer is activated (see INITIALIZATION section of manual).

HINTS ON CONNECTION:

Information on the input signal levels required are in the main manual. Your audio input ground (shield) should go to the "sleeve" or outer conductor of the phone plug. The "hot" lead (usually red in color) should be connected to the "tip" or center conductor of the phone plug. If you use a stereo plug, the "ring" conductor is connected to the analyzer board but does nothing. Strong signals on this wire may cause "crosstalk" and should be avoided.

If, during use, you observe that the bar corresponding to 16kHz seems unduly high, or present with no input, this is almost certainly due to radiation from the TRS-80 monitor. We have found that this happens frequently, especially when the analyzer gain is fairly high. Possible cures include careful shielding of the audio or CRT cables, physically moving the CRT, or, if necessary, providing more external amplification or even turning off the CRT when a measurement is being made.

APPENDIX #2: PERFORMANCE ASSURANCE, TEST AND CALIBRATION

The purpose of this section is to enable the user to assure himself/herself that the Analyzer is functioning properly. There are two levels of testing desirable. The first is to create assurance that the unit is not defective in some serious fashion, such as having malfunctioning conversion circuitry, filters, or power supply components. The second is to ascertain that the unit is properly calibrated and within specifications as to frequency and amplitude. If both sets of tests are completed successfully, the chances are excellent that the Analyzer is not defective in any fashion.

GROSS DEFECT TESTS

ROM CHECKSUM

This test assures that the memory mapping circuitry, machine language ROM, and to a lesser extent, the power supply circuitry is functioning properly. Enter the following into the TRS-80.

```
100 REM ---COMPUTE EPROM CHECKSUM---
140 OUT 231,144
160 OUT 229,0
180 CLS: PRINT "COMPUTING CHECKSUM OF ROM DATA"
190 FOR X%=-1024 TO -1
200 C=C+PEEK(X%)
210 NEXT X%
220 CH=C-(INT(C/256))*256
230 CH=(CH/16)-14: PRINT "CHECKSUM=";CH
```

The checksum should be zero. If it is not, it means that in all likelihood there is an error in the ROM. This may be verified by repeating the test several times. If the answers are identical, it is the ROM. If they change, it may be the power supply.

AUDIO TEST

CLOAD the INTERACTIVE operation program on the tape cassette furnished with the unit and connect the output of a sine wave audio generator to the Analyzer input connector. The generator level should be between +14 and -16dBm and the output frequency should be 1kHz. RUN the program and press the +6dB and +3dB keys as necessary until the highest bar on the graph is near full scale. Place the cursor under this bar and confirm that the frequency display reads "1KHZ". Also confirm that the bars on either side of the 1KHz bar are shorter and monotonically decreasing on either side. Now, manually sweep the oscillator from 20Hz to 20KHz and confirm that, although the position of the maximum bar varies, the display doesn't do anything strange, such as lose a bar, or have one suddenly stick up at the wrong end. Also confirm the the LEVL bar remains at an essentially constant amplitude throughout this procedure. (If a low level input is being used, the LEVL bar may drop by a dB or so at the high-frequency end). Note that as the frequency varies between bars it is normal for both adjacent bars to lose some height. Increase and decrease the Analyzer gain with the 3dB control keys and observe that the gain varies smoothly between steps. Note that it will be necessary to vary the signal generator level during this step so the the main bar does not exceed full scale. If it does so by a significant margin, the analyzer pre-amplifier can "clip" and generate serious harmonic distortion which will appear as certain bars getting too high out of "order".

If all of the above seems OK, the analyzer is probably working properly. If any anomalies appear, they should most certainly be found and corrected before any calibration is attempted.

CALIBRATION

To calibrate the Analyzer, you will need an audio signal source capable of generating sine waves at all frequencies between 20Hz and 20kHz, at a constant level. Most Wien bridge oscillators and "function generators" meet these criteria. If your generator does not supply a constant level, then an ac voltmeter may be used to reset the level after changing frequency. A frequency meter connected to the output of the oscillator will be helpful in determining whether the filter center frequencies are within specification.

There are 31 potentiometers on the filter circuit board for adjusting the GAIN of the individual filters, and one on the interface board for adjusting offset voltage. The GAIN potentiometers have virtually no effect on the filter center frequency and may not be used to adjust it. It is VERY important to adjust the unit using the procedure shown, and NOT to set the oscillator to the theoretical center frequencies and then adjust the GAIN pots for a specific level! Note that there is no adjustment for the "LEVL" amplifier. This output serves as a reference for all the other settings.

Place two pieces of translucent tape (Scotch Magic (TM) transparent tape works OK) across your monitor screen, one whose bottom is slightly below the top graduation of the display, and another whose top is about 1 division below the bottom of the first. The purpose of this is to establish a reference line on the screen which is easily identifiable. Run the INTERACTIVE program and set the input level to at least 0dBm. Use the +3dB and -3dB keys with the display mode in "LIN" to bring the "LEVL" bar to the clear line between the two pieces of tape. It will probably be necessary to trim the generator output slightly to get the height exactly right. This will be the level reference for the rest of the procedure.

Temporarily remove the generator input.

Open the chassis and adjust the OFFSET pot so that all the filter bars are just turned off. It is normal for the LEVL cell to be slightly higher than the filter cells with no input signal. Now adjust the oscillator to near 20Hz and carefully tune it so that the 20Hz bar is at its maximum amplitude. Adjust the 20Hz GAIN pot on the filter board so that the level is within the tape slit and as nearly equal to the LEVL bar as possible. Again tuning the oscillator carefully to the maximum bar height, adjust the 25Hz gain pot for the same bar height. If you wish, note the frequency counter readings for each adjustment to calculate the center frequency tolerance. Continue this procedure for all 31 GAIN controls. The parts location drawing will give the order in which they are to be adjusted.

This completes calibration of the Spectrum Analyzer.

APPENDIX #3: PINK NOISE GENERATOR

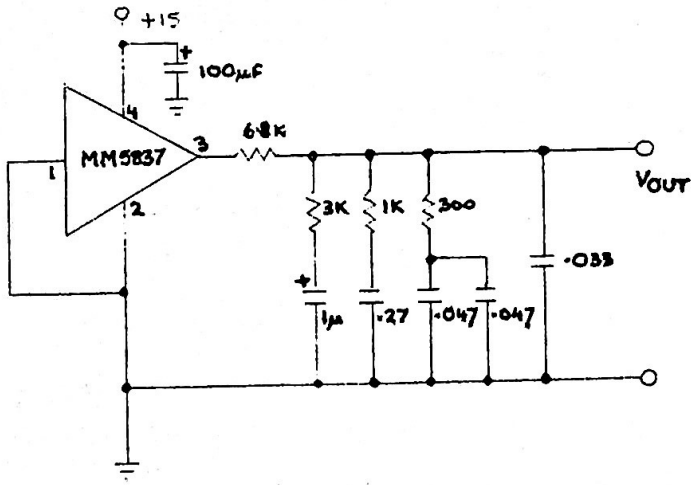

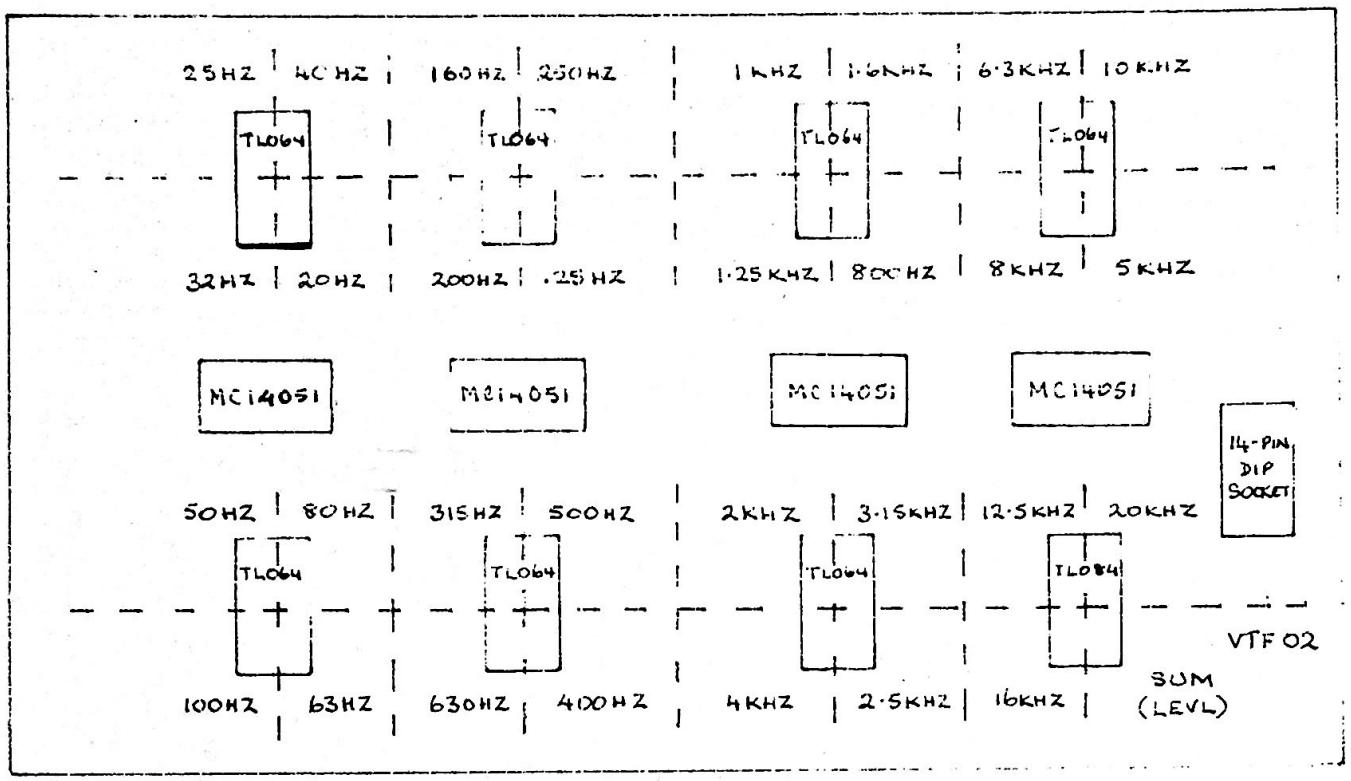

VTF02 FILTER BOARD PARTS LOCATION DIAGRAM

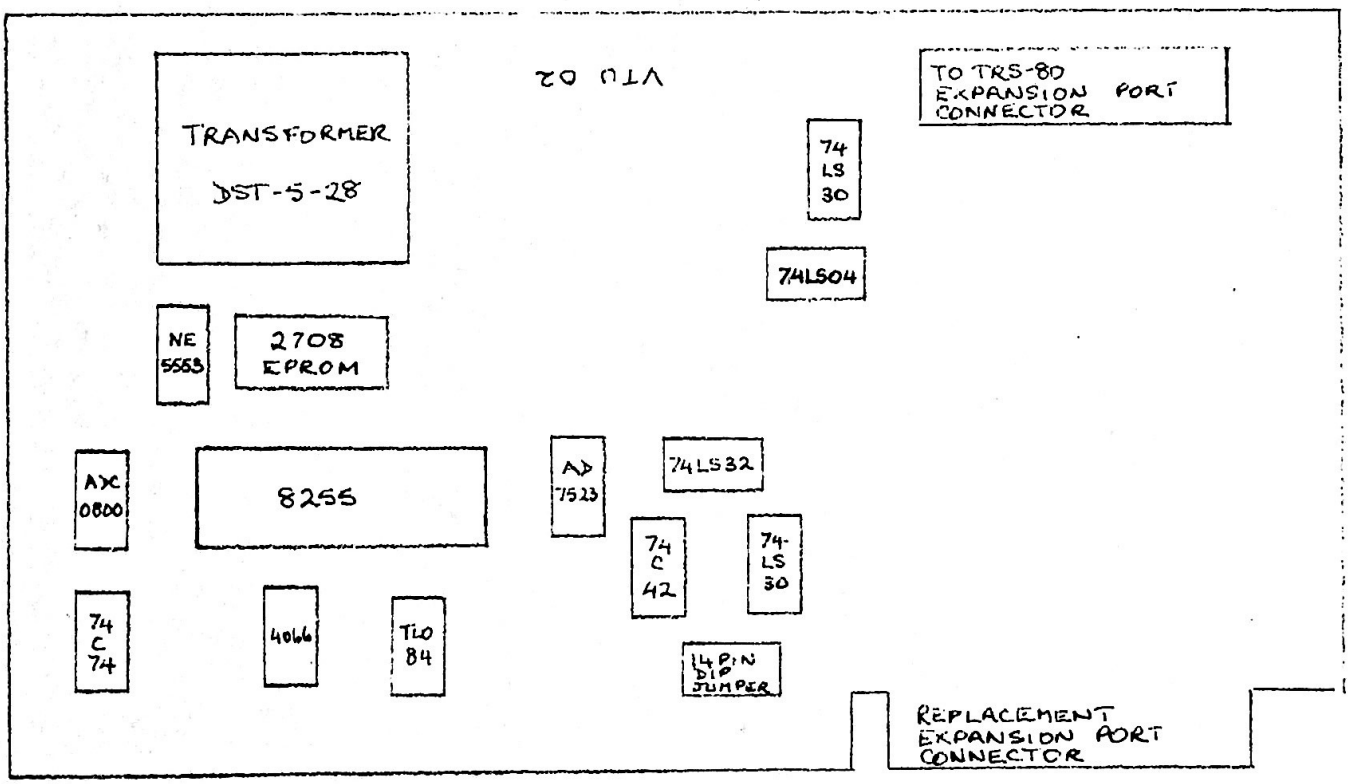

VTU 02 CONTROL BOARD PARTS LOCATION DIAGRAM

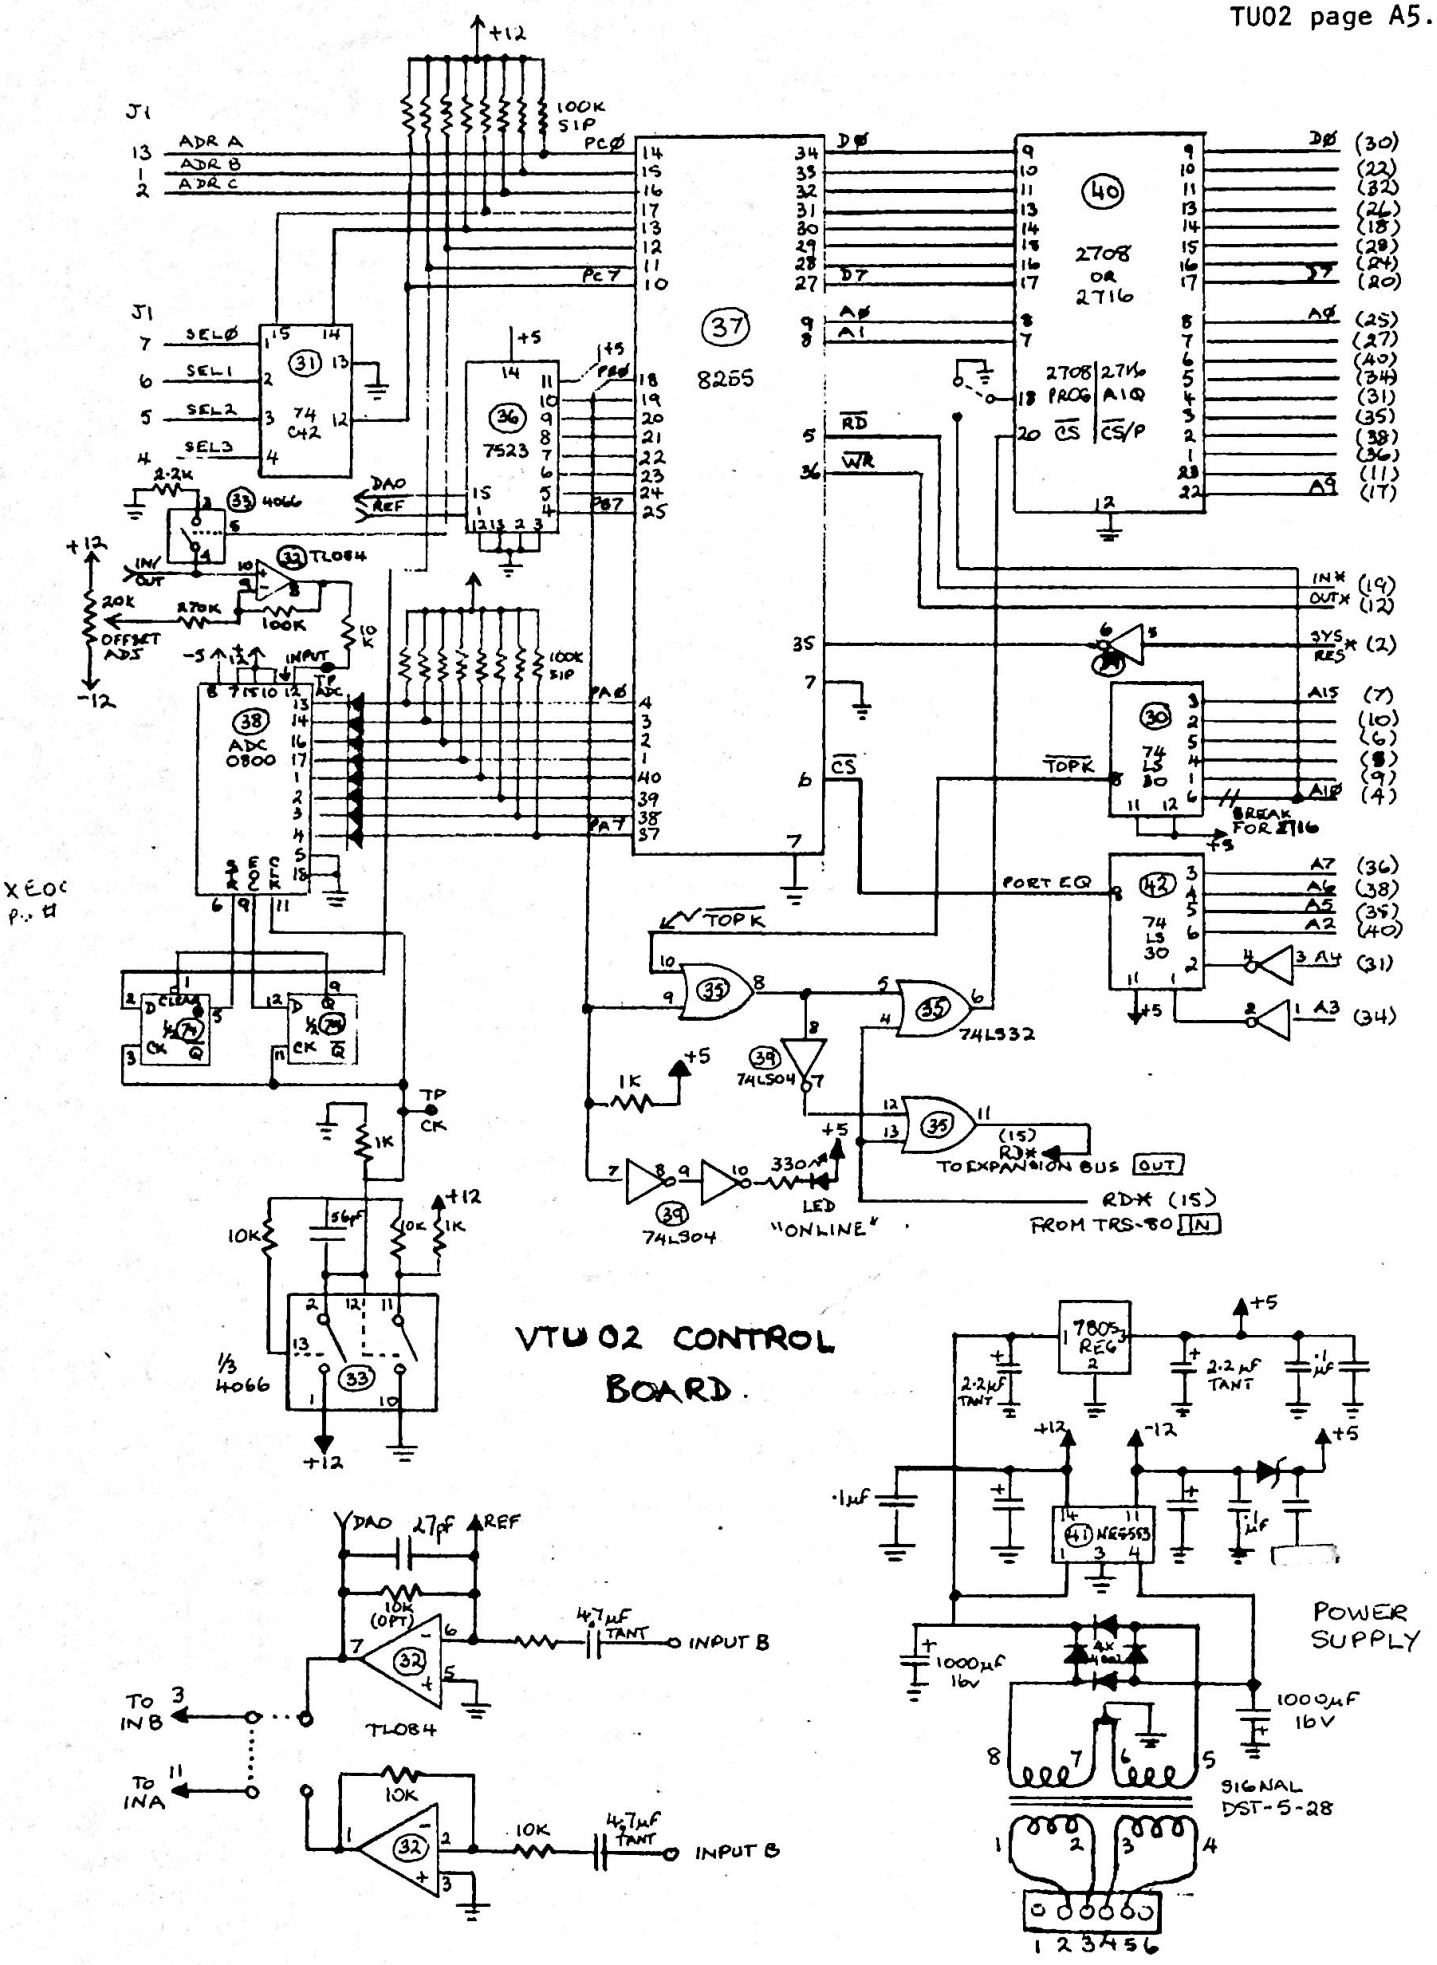
VTU02 CONTROL BOARD.
J1
ADR A
ADR B
ADR C
SEL0
SEL1
SEL2
SEL3
74 C42
4066
TL084
OFFSET ADJ
ADC 0800
8255
7523
2708 OR 2716
74 LS 30
74LS32
74LS04
LED "ONLINE"
TO EXPANSION BUS OUT
RD* (15) FROM TRS-80 IN
BREAK FOR 2716
PORT EQ
TOPK
XEOC
1/3 4066
INPUT B
TO INB
TO INA
7805 REG
NE555
POWER SUPPLY
SIGNAL DST-5-28

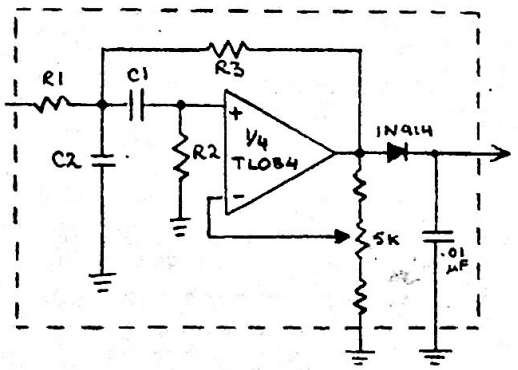

TYPICAL 1/3 OCTAVE FILTER SECTION

VTF02 FILTER BOARD.

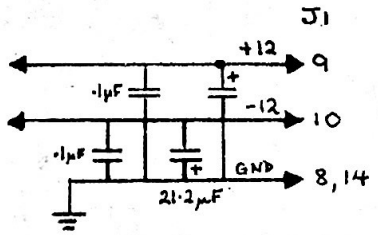

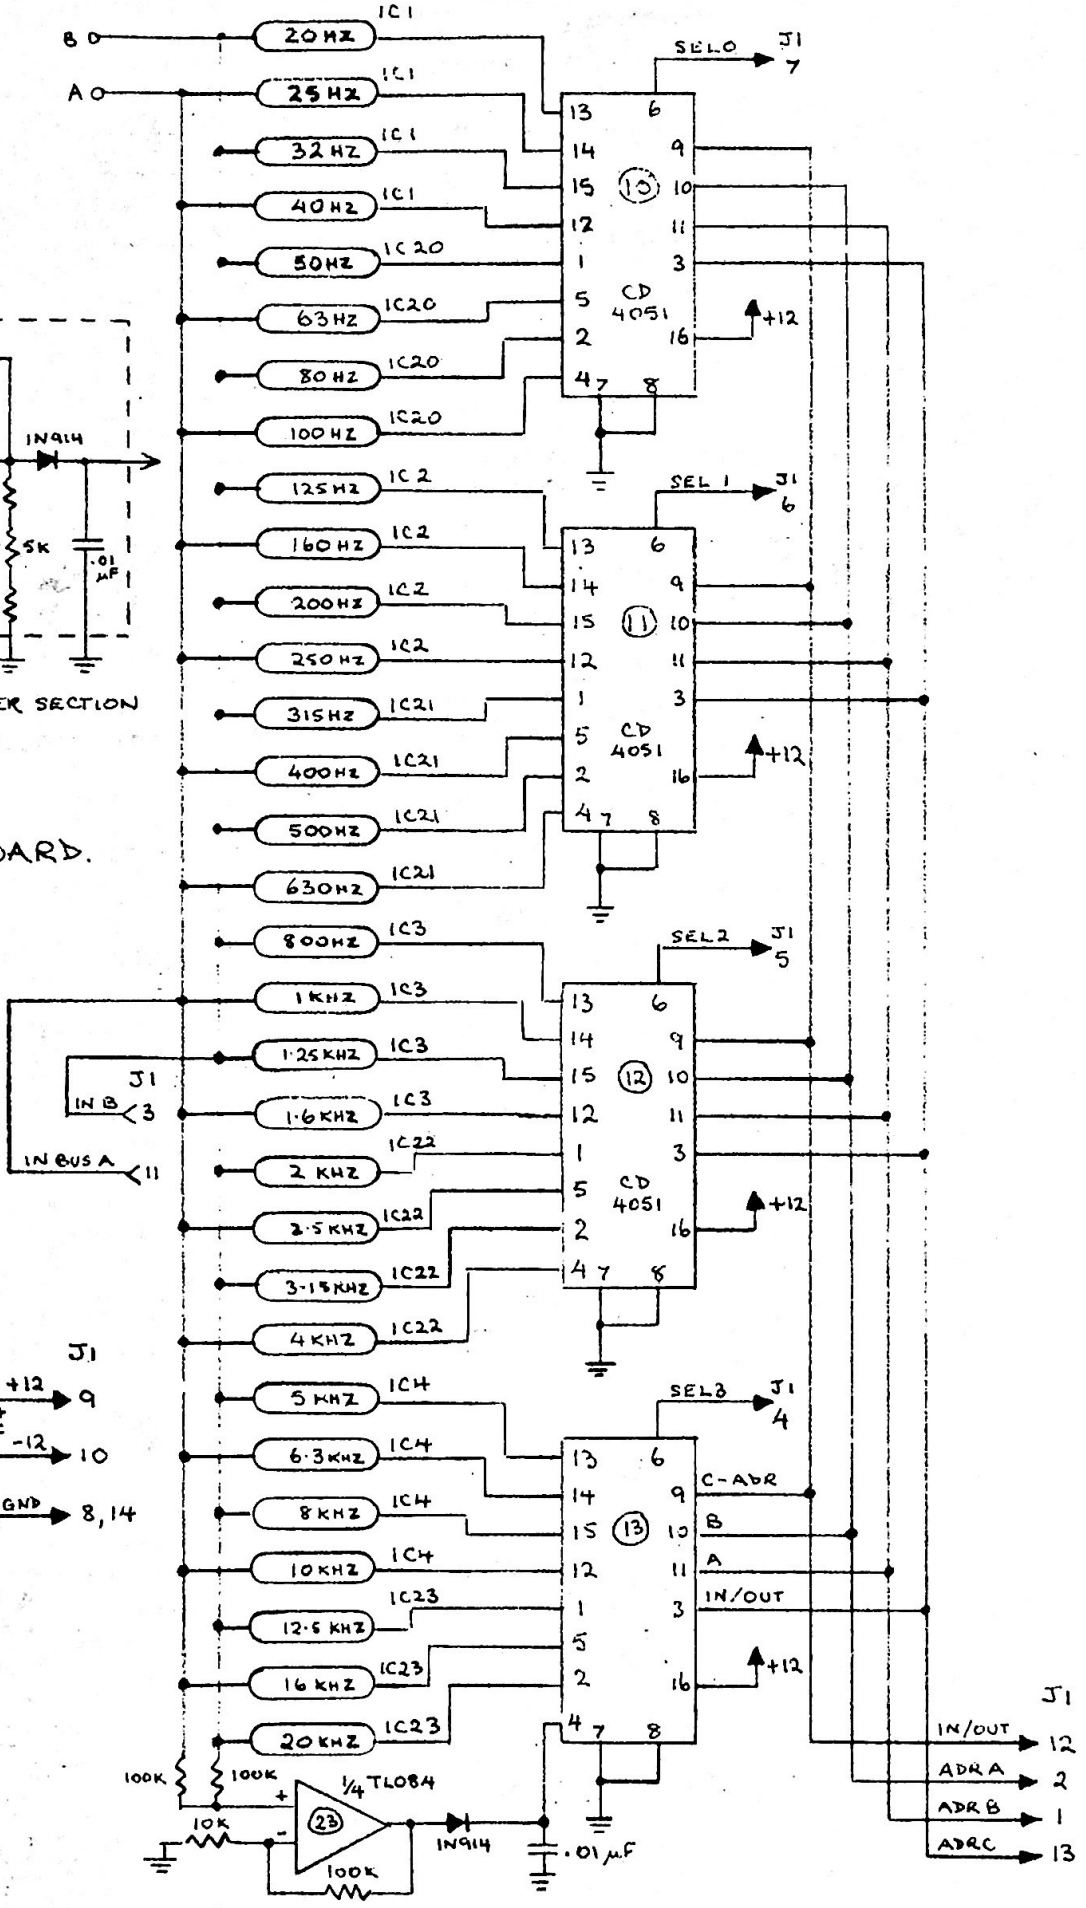

APPENDIX #4: INSTRUCTIONS FOR INTERACT PROGRAM

CLOAD PROGRAM
COMMAND "RUN"

The keys of the TRS-80 will have the following functions:

↓	GAIN decrease by 3 dB	Repeating if held
↑	GAIN increase by 3 dB	Repeating if held
1	Alternates LOG/LIN	Repeating if held. Displays 'LOG 1DB/DIV' or 'LIN 120MV/DIV' in uppper left corner
2	Alternates AVERAGE/SINGLE	Repeating if held. In AVERAGE mode, displays 'AVERAGING' under LOG/LIN display
3	PEAK HOLD	Displays 'PEAK HOLD' or 'DECAY' between 'LIN/LOG' and 'AVERAGING'
4	VERY SLOW DECAY	
5	SLOW DECAY	
6	MED DECAY	
7	FAST DECAY	
← →	Move Frequency Cursor	Repeating if held. Cursor frequency displayed in bottom left corner

CRT LAYOUT

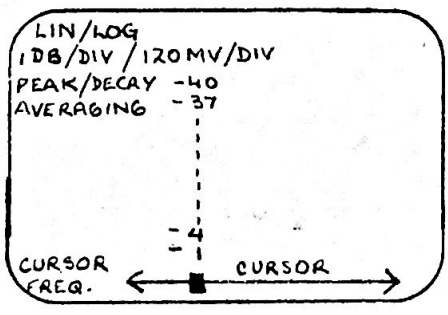

(NOT TO SCALE)

KEYBOARD LAYOUT

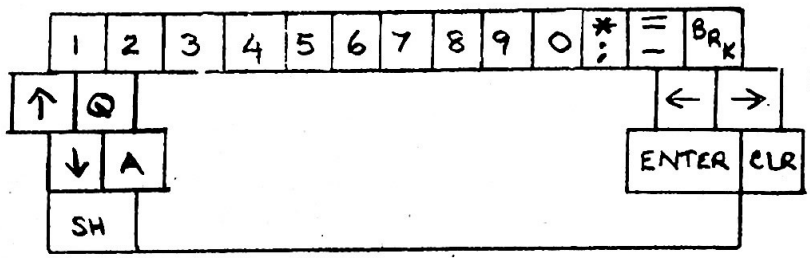

Eventide
the next step

digital delay lines
instant flangers
omnipressors
harmonizers
monstermats

EVENTIDE CLOCKWORKS, INC. • 265 WEST 54TH STREET • NEW YORK, N.Y. 10019 • 212-581-9290

TRS-80 is a trademark of Radio Shack/Tandy Corporation
Harmonizer, Instant Flanger, Monstermat and Omnipressor are trademarks of Eventide Clockworks Inc.